

FIRST DATA FROM THE SECOND ITERATION
OF THE DARKSIDE 10 KG DETECTOR

Alden Fan

Adviser: Peter Meyers

Second Reader: Frank Calaprice

PRINCETON UNIVERSITY

May 2, 2011

Pledge

This paper represents my own work in accordance with university regulations.

Alden Fan

Acknowledgements

First and foremost, I thank Professor Peter Meyers for giving me the chance to participate in the DarkSide experiment. His consummate care and guidance throughout this senior thesis has taught me most of what I know about experimental physics. Professor Meyers's dedication to his work and his students has inspired me to pursue physics beyond the undergraduate level.

I would also like to thank Ben Loer, who wrote much of the code for DarkSide, on which my analysis was built; he helped me on numerous occasions to navigate the murky waters of ROOT. And I thank Jason Brodsky, who wrote the photoelectron calibration code, on which some of my analysis relies and which I used as a template for much of my own code.

I am eternally grateful to all of my friends and family, who have encouraged me, comforted me, and believed in me. I especially acknowledge my seven "roommates" for making the last four years a spectacular adventure. Finally, I am indebted to Rebecca Yu, whose vivacity inspired me when I needed it most.

Abstract

First Data from the Second Iteration of the DarkSide 10 kg Detector

Alden Fan

Dark matter is known to dominate the matter content of the universe, and its identification is one of the main challenges of modern physics. One of the most promising methods is to detect the scintillation light produced by the interaction of a dark matter particle with a target material such as liquid argon as is being done in the DarkSide experiment. This experiment has completed its second iteration of a 10 kg prototype detector. The main challenge of the experiment is to detect and identify the large background produced by neutrons and beta and gamma rays, which produce different scintillation pulse signatures. One of the main methods of background rejection is to use the ratio of the fast and slow component of the scintillation pulse to discriminate between nuclear and electronic interactions. We present here a description of the first data coming from the second iteration of the DarkSide detector and a detailed analysis of the slow component lifetime in the primary scintillation pulses. Using a χ^2 minimization, we fit an exponential function to a summed waveform with corresponding errors, which we motivate by a photoelectron counting statistics argument. We discuss six cuts on the data to exclude events with cosmic ray muons, afterpulses, false triggers, and other anomalies. We find that the estimated slow component lifetime is dependent on multiple parameters, so an absolute value cannot be given at this time, but that it is uncorrelated with variations in the impurity levels of this iteration of the DarkSide detector.

Contents

Introduction	1
1 Dark Matter	3
1.1 Historical evidence	4
1.2 Candidates	6
1.3 Search Methods	7
1.4 Current indirect detection experiments	8
1.5 Direct Detection Experiments	11
2 LAr TPC Experiments and DarkSide	20
2.1 Liquid argon	20
2.2 Scintillation detection	22
2.3 Ionization detection	23
2.4 Photomultiplier tubes	24
2.5 Active veto	26
3 Data and Cuts	27
3.1 ^{22}Na coincidence runs	27
3.2 Data output and analysis modules	28
3.3 Single photoelectron calibration	33
3.4 Slow-component lifetime	33
3.5 Baselines	35
3.6 Cuts	36

4	Photoelectron Counting Statistics	41
4.1	Summed waveform	42
4.2	Alternating behavior	44
4.3	Uncertainties	44
5	Fit and Results	52
5.1	Fit	53
5.2	Anomalous structure	53
5.3	Fit window	57
5.4	Pre-trigger vs. post-pulse	59
5.5	Lifetime over time	62
6	Conclusion	65
A	Sinusoidal Fluctuations	67
	References	73

Introduction

General relativity has successfully described our universe on the largest of length scales, while the Standard Model has successfully described the universe on the smallest of length scales. But both these models leave many questions open. Why is there no unification of the forces? Why are the mass scales the way they are? Why is the expansion of the universe accelerating?

The Standard Cosmological Model is currently the most widely accepted description of how the universe was formed and predicts the mass-energy content of the universe [1]. Recent measurements from the Wilkinson Microwave Anisotropy Probe (WMAP) place the tightest limits on the model, stating that the universe is composed of 73% dark energy, 23% nonbaryonic dark matter and 4% baryonic matter [2]. Dark energy is used to account for the acceleration of the expansion of the universe, and dark matter is used to account for the observed motion of stars and galaxies [1, 3]. Though these entities are predicted to exist, we know essentially nothing about them.

In this thesis, we give a broad overview of dark matter and prospects for its detection as well as take a closer look at direct detection efforts, particularly the DarkSide experiment currently under development at Princeton University. We then present an analysis of some of the first data from the second iteration of the 10kg prototype version of that experiment, focusing on the slow-component lifetime, a crucial piece of the detection method. In Chapter 1, we give an overview of the historical evidence, candidates, and search methods for dark matter. In Chapter 1.5 we go into more detail about direct detection experiments, including their main challenges, and the current status of the major experiments in the field. In Chapter 2 we give a more detailed description of Time Projection Chamber (TPC) and Liquid Argon

(LAr) technologies and their usefulness in dark matter detection. We also highlight some of the unique features of the DarkSide experiment. In Chapter 3, we begin the detailed analysis of the slow-component lifetime, discussing the data and cuts on the data that we used. In Chapter 4, we discuss the photoelectron counting statistics that motivated summed waveform and error calculation that we derive. Finally, in Chapter 5, we present the results of the analysis, including several anomalies in the summed waveform and the behavior of the slow-component lifetime over time for the DarkSide 10 kg detector.

Chapter 1

Dark Matter

The dark matter problem of the 20th and 21st century physics is conceptually akin to the 19th century challenges of discovering other planets in the Solar System [4]. Newton’s laws of gravitation were widely accepted as the governing principle for describing the motion of astrophysical bodies, and any body’s deviation from its expected path led to one of two conclusions: either they indicated the presence of an as yet unseen object, or they were regarded as a refutation of the current laws of physics. For example, in 1846 irregularities in the path of Uranus led to the discovery of Neptune [5], while discrepancies between the observed and theoretical orbits of Mercury were not explained until Einstein discovered the general theory of relativity [6]. Similarly, today there are large discrepancies between the observed and predicted motions of stars and galaxies, which are only explained by assuming the existence of a large amount of an as yet unseen form of matter, or by assuming a departure from the current theory of general relativity. The general consensus is to assume an unknown particle, dubbed “dark matter.” We give here a brief history of the evidence for dark matter, followed by an overview of the candidates and detection methods for it, including a survey of the major experimental efforts in the field.

1.1 Historical evidence

The idea that visible matter constitutes only a small portion of all the mass in the universe began with the measurement by Fritz Zwicky in 1933 of eight galaxies within the Coma cluster with unexpectedly high radial velocities, leading him to estimate a mass density 400 times greater than that expected from luminous matter [7]. Zwicky was one of the first to understand that the missing mass from the Coma cluster, and possibly other galaxy clusters, represented a significant unsolved problem in astrophysics [8].

In 1970 Vera Rubin found more evidence that there is much more matter in the universe than expected by measuring the rotation curves¹ for stars and gas clouds within many galaxies [9]. From Newtonian dynamics, the expected circular velocity of stars and gases about a galactic center should follow Newton's law of gravitation,

$$F = G \frac{mM(r)}{r^2} \quad (1.1)$$

where m is the mass of the rotating body, r is the radial distance from the galactic center to the body, and

$$M(r) = 4\pi \int_0^r \rho(r') r'^2 dr' \quad (1.2)$$

where $\rho(r)$ is the mass density profile of the galaxy. Assuming circular motion of the stars and gases, such that $F = mv^2/r$, then the circular velocity of the body should follow

$$v(r) = \sqrt{\frac{GM(r)}{r}}. \quad (1.3)$$

Beyond the optical disc, where the bulk of the luminous matter of the galaxy resides [3], $v(r)$ should have a $1/\sqrt{r}$ dependence. But Rubin found that rotation curves were constant well beyond the optical disc for many galaxies [9, 10], indicating the existence of a mass halo with $M(r) \propto r$ and $\rho(r) \propto 1/r^2$ [4]. Figure 1.1 shows a constant rotation curve for such a galaxy along with the gas, disc, and dark matter halo contributions.

¹A rotation curve relates the velocity of a body around an orbit to its distance from the center.

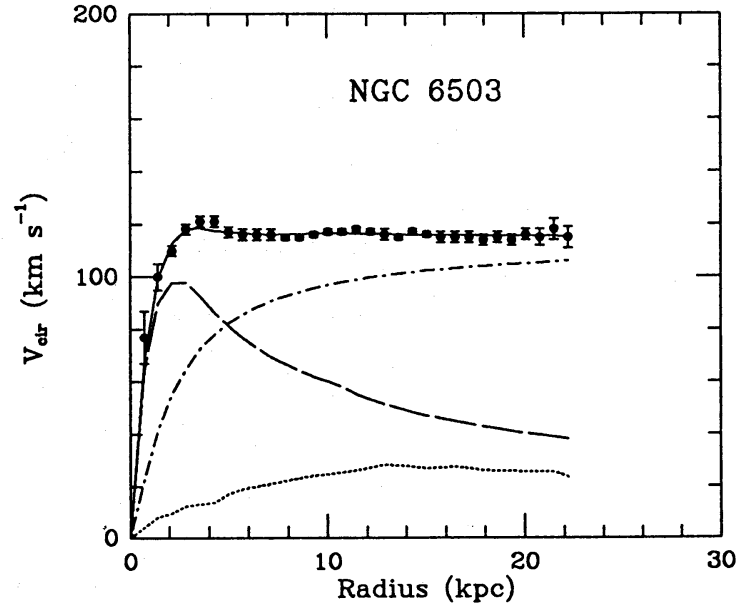


Figure 1.1: The rotation curve for the NGC 6503 galaxy. The dotted, dashed, and dashed-dotted lines are the contributions of gas, disc, and dark matter, respectively, to the total rotation curve. Image taken from Ref. [11].

By the mid 1970s, the majority of astronomers accepted the fact that significant amounts of missing mass existed in all galaxies [8]. However, the particular form of the mass was yet unknown; possibilities included brown dwarfs, white dwarfs, black holes, very hot gas, or some as yet unsuspected form. Not until the 1980s did our more modern understanding of the universe come about, in which cold dark matter dominates the matter content and baryonic matter constitutes only a small percentage.

More recent measurements using gravitational lensing provide further evidence for the existence of dark matter. As light passes by a large mass distribution, its path is bent through space-time, which can be used to determine the amount of mass present. This has been used to study, for example, the Bullet Cluster, in which a subcluster is passing through a larger cluster of galaxies [12]. They find that the resulting motion of the galaxies can only be explained if there is a large amount of unobserved matter present.

1.2 Candidates

There are dozens of models that attempt to explain the dark matter content of the universe, leading to a zoo of particle dark matter candidates. A successful dark matter candidate must satisfy certain properties [13]. First, it must be gravitationally interacting to explain the observed motion of stars and galaxies within a cluster. It must not be electromagnetically or strongly interacting; otherwise, we would have detected it long ago. Note that it may interact via the weak force or by some other as yet undiscovered force on the weak scale. The dark matter particle must be stable, since it is expected to have been present since the early universe. And it must be cold, meaning that dark matter should have been non-relativistic near the beginning of the formation of galaxies [3]. We discuss here the most popular ones.

The most popular candidate is the Weakly Interacting Massive Particle (WIMP), which encompasses a class of particles that satisfy the above properties. Theoretically, WIMP candidates arise from Supersymmetry, which is a widely discussed extension of the Standard Model of particle physics. The most popular Supersymmetric dark matter candidate is the neutralino, though there are other Supersymmetric candidates such as the gravitino and the Kaluza-Klein particle [14]. Though Supersymmetry is based on a solid theoretical foundation, there is no experimental evidence for it. In fact, one of the main goals of the CMS and ATLAS experiments at the Large Hadron Collider in Geneva, Switzerland is to find such evidence.

WIMPs are particularly attractive because, making reasonable assumptions about its mass ($\sim 100 \text{ GeV}/c^2$; the proton has mass $\sim 1 \text{ GeV}/c^2$) and interaction cross-section close to that of weak interactions ($\sim 1 \text{ pb}$), then the abundance of dark matter very nearly matches that expected for thermal relics from the Big Bang [16]. This is referred to as the “WIMP miracle.”

The other popular dark matter candidate is the axion, which is motivated as a solution to the strong CP problem [15]. If discovered, it would be the lightest of elementary particles with nonzero mass. The Axion Dark Matter Experiment (ADMX) is currently biggest effort looking for axions.

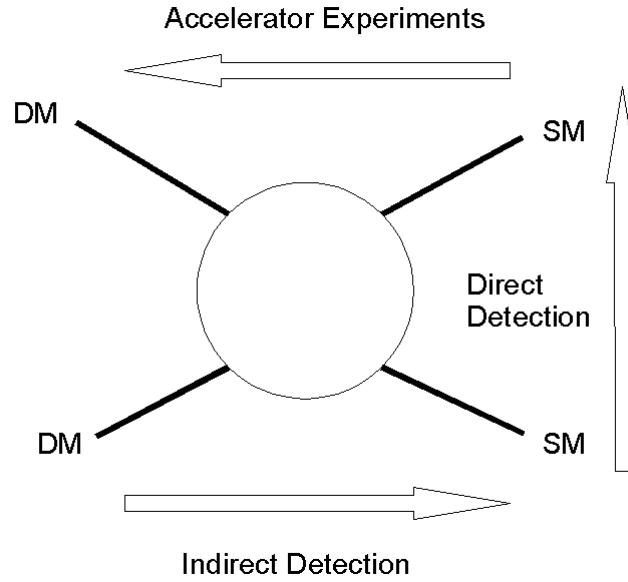


Figure 1.2: Three different approaches to detecting WIMPs. DM denotes a dark matter particle and SM denotes a Standard Model particle [16].

1.3 Search Methods

There are three main approaches to dark matter detection: indirect detection, direct detection, and detection using accelerators, which are depicted schematically in Fig. 1.2.

In indirect detection, we assume that dark matter particles annihilate into Standard Model particles, which we can detect as an excess of energy in a detector. The difficulty is that one must be able to account for all other particles observed due to known processes, which are many [16].

Accelerators take the opposite approach to indirect detection. Known Standard Model particles are smashed together in accelerators such as the LHC in the hopes that they will annihilate and produce some unknown form of matter. Whether these particles are dark matter is then an additional step to be determined separately by some other form of detection. Even if they are not such particles themselves, the new particles would at least lead toward a model that includes some form of Dark Matter [16]. If the annihilation does produce a dark matter particle, it will be very difficult

to detect and its stability (on a cosmological timescale) will be impossible to prove [17]. Its presence would most likely only be in the form of missing energy.

Lastly, in direct detection experiments, one looks for events in which a dark matter particle interacts with a Standard Model particle. Near the Earth, 100 GeV/ c^2 WIMPs have an estimated density of ~ 0.3 GeV/ c^2 /cm³ and an estimated local velocity of 270 km/s [18], corresponding to a local flux of $\sim 10^5$ cm⁻² s⁻¹. The WIMP is expected to scatter elastically off a known target nucleus, transferring some of its kinetic energy to the nucleus, which is then detectable. However, the WIMP interaction cross section is predicted to be extremely small; that is, it has a very small probability of interacting with anything. Since the interactions are expected to leave very small energies and occur exceedingly rarely, the detectors must be extremely sensitive. Thus there will be a very large background from known cosmogenic and terrestrial sources. The primary challenge of direct detection experiments is then to reject all of this background. Thus direct detection experiments are often likened to looking for a needle in a haystack.

1.4 Current indirect detection experiments

We highlight the major experiment results from recent indirect detection experiments. These experiments are generally cosmic-ray detectors sensitive to charged particles originating from a distribution of distant sources mainly associated with supernova remnants and pulsars, which are highly magnetized rotating neutron stars [19, 20].

1.4.1 ATIC

The Advanced Thin Ionization Calorimeter (ATIC) instrument is a balloon-borne detector designed to measure the composition of cosmic rays, particularly electrons in the energy range of 300 to 800 GeV [21]. The ATIC instrument's results, which took data in 2002-03, are shown by the solid red dots in Fig. 1.3, along with data from previous measurements. The solid curve is based on simulations using GALPROP, a cosmic ray propagation package, of the electron content of cosmic rays [22]. Below

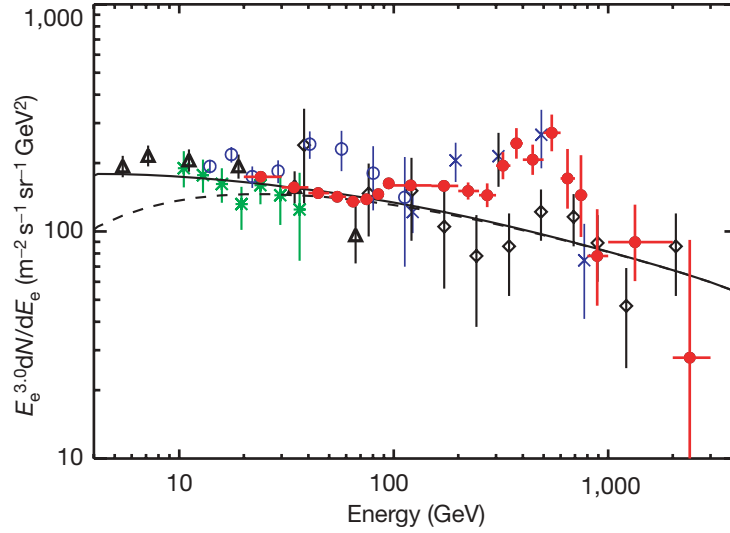


Figure 1.3: The ATIC results (red) along with previous measurements and the predicted curve using GALPROP simulations (see Ref. [21] and references within).

100 GeV the ATIC results match both the theory and previous measurements quite well. However, between 100 GeV and 650 GeV, there is an excess of electrons, and then the spectrum drops rapidly. One explanation for the structure attributes it to the annihilation of dark matter, which according to some models produces electron-positron pairs [23]. The spectrum featured in the ATIC data would correspond to a WIMP mass of 620 GeV [21]. However, the data is also consistent with electrons being produced by a nearby astrophysical source, such as a pulsar or a micro-quasar, which is an energetic binary system of a normal star paired with a neutron star or black hole [21, 24]. However, this would be the first direct observation of a nearby source of particles with energies in the hundreds of GeV. Thus, though the ATIC data cannot conclusively point to the observation of dark matter annihilation, it opens up new windows for study with either explanation.

1.4.2 PAMELA

The Payload for Antimatter and Matter Exploration and Light-nuclei Astrophysics (PAMELA) experiment is a satellite observatory designed to detect charged particles

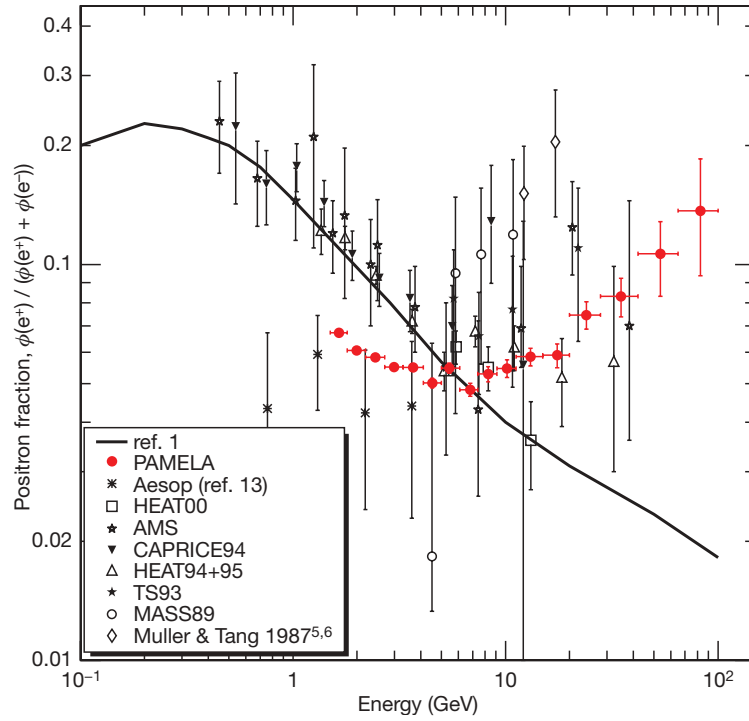


Figure 1.4: The positron fraction from the PAMELA experiment (solid red dots) as well as previous measurements and the predicted fraction based on secondary processes (see Ref. [25] and references within).

within cosmic rays, specifically the small fraction of positrons and antiprotons with energies of 1.5 to 100 GeV [25]. In particular they are interested in measuring the ratio of the positron flux to the sum of the electron and positron fluxes,

$$\frac{\phi(e^+)}{\phi(e^-) + \phi(e^+)} \quad (1.4)$$

which is called the positron fraction. Because positrons are believed to be mainly created by the interaction of cosmic ray nuclei with the interstellar medium [26], referred to as a secondary process, this measurement will lead to a better understanding of the nature and distribution of particle sources in our Galaxy. The results from a data run from July 2006 to February 2008 are shown in Fig. 1.4, where the solid line is plotted from calculations based on the secondary processes. The PAMELA results do not match up very well with previous results nor the calculation. In particular, at energies above 10 GeV, the positron fraction increases significantly with increasing

energy. The PAMELA collaboration attributes this rise to positrons from primary sources, as opposed to positrons from secondary processes [25]. One tantalizing primary source is from the annihilation of dark matter. However, WIMP particles are predicted to annihilate into both high energy antiprotons and positrons [4], yet the PAMELA collaboration previously measured the antiproton-to-proton flux ratio to match that predicted from secondary production calculations for antiprotons [27]. This means that if the rise in the positron fraction is to be due to WIMPs, they must annihilate into mostly leptonic states. Finally, the increasing positron fraction is, like the ATIC data, compatible with Standard Model processes from astrophysical primary sources such as pulsars and microquasars [25]. Thus PAMELA cannot conclusively claim to have observed dark matter annihilation.

1.4.3 Fermi-LAT

The Large Area Telescope (LAT), the main instrument aboard the Fermi Gamma-Ray Space Telescope mission, was designed primarily to survey the gamma-ray sky between 20 MeV and 300 GeV [28]. It is also a detector for cosmic ray electrons and positrons, though it cannot differentiate between them. Based on data collected from August 2008 through January 2009, the Fermi-LAT collaboration finds an excess of electrons and positrons in the energy range above about 100 GeV, as can be seen in Fig. 1.5. Also plotted are previous measurements of the cosmic ray electron and positron spectrum and the predicted curve based on a GALPROP calculation. In agreement with the ATIC results, Fermi-LAT finds an excess of electrons and positrons in the hundreds of GeV range. Thus many experiments are on the verge of discovering new physics. But whether it is dark matter or an undiscovered astrophysical body has yet to be determined.

1.5 Direct Detection Experiments

The goal of direct detection experiments is to detect and identify a WIMP by observing its interaction with a known material. The WIMP will deposit some of its kinetic

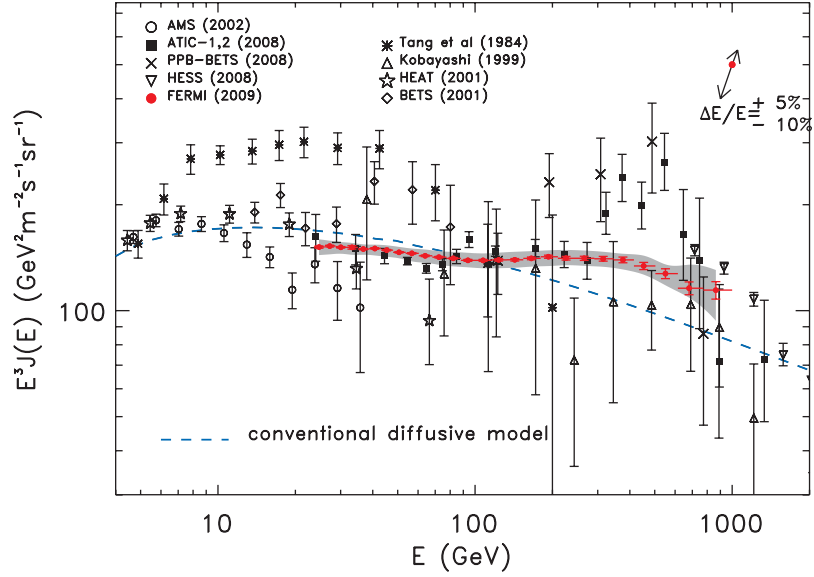


Figure 1.5: The results from the Fermi-LAT with previous measurements of the cosmic ray electron spectrum and the predicted from a GALPROP calculation (see Ref. [28] and references within).

energy by elastically scattering off a target atomic nucleus. We then try to detect the recoil energy of the nucleus and from that identify it as a WIMP. The difficulty is that the detectors are also sensitive to energy deposits from many other particles. Furthermore, since the WIMP interaction cross section is predicted to be exceedingly small, meaning very few interactions are expected to occur within the detector, we must be able to reject the background extremely efficiently. There are two main types of background in direct detection experiments: electron and gamma backgrounds and neutron backgrounds.

1.5.1 Backgrounds

Neutron backgrounds

Neutrons are the most dangerous form of background because they can produce signals identical to that expected from WIMPs. Thus it is very important to effectively shield the detector from neutrons. Neutrons may enter the detector from any number of sources, including cosmic ray muons interacting near the detector, environmental

radiation, and radioactive contamination of materials in the detector. Therefore direct detection experiments are located deep underground to shield against cosmic ray muons, are heavily shielded using lead or other materials, and contain materials with as low radioactivity as possible. The only method for discriminating between neutrons and WIMPs makes use of the fact that, if the detector is large enough, neutrons are much more likely to scatter more than once within a detector. Therefore there is great advantage in being able to scale a detector up to large volumes as it increases the probability that a neutron will scatter more than once.

Electron and gamma backgrounds

The most abundant background is from electrons and photons. Unlike neutrons and WIMPs, which interact with the target nucleus, electrons and photons interact via the photoelectric effect and Compton scattering. Since different interaction mechanisms produce different energy signatures, then if one is able to measure the total energy of an interaction, it is possible to deduce whether it was from an electron or gamma or from a neutron or WIMP, providing an effective rejection method for the former. Experiments employ several different detection techniques that allow for such background rejection: scintillation, ionization, or phonon detection, or any combination thereof. Scintillation is the light produced by certain materials in the presence of ionizing radiation (see Sec. 2.2 for a more detailed discussion). The ionization channel refers to the detection of disassociated electrons from their nuclei due to interactions with an energetic particle. Phonon detection refers to the detection of an increase of heat in a target material due to some form of energy deposition. There are various technologies for detecting scintillation, ionization, and phonons, and those that can detect interactions through multiple channels are particularly attractive since they allow for the most efficient background rejection. For example, crystals at cryogenic temperatures (in the mK range) are typically used for phonon and phonon/ionization detection, while liquid noble gases are used for scintillation and scintillation/ionization detection.

1.5.2 Current direct detection experiments

There are currently dozens of direct detection experiments in various states of development from R&D to data-taking. We describe in the rest of the chapter a small selection of experiments that use a variety of detection techniques. See Ref. [29] for a more extensive list of WIMP-search experiments.

CDMS II

The Cryogenic Dark Matter Search (CDMS II) experiment, located at the Soudan Underground Laboratory, uses a series of 19 Ge (4.75 kg total) and 11 Si (1.1 kg total) particle detectors at cryogenic temperatures (<50 mK) in a combination of phonon and ionization detection [30]. An electric field is applied across each detector for the ionization detection, and the ratio of the energies measured from ionization versus phonons provides the primary electronic background rejection. However, electronic events near the surface of the detectors can fail this cut, so an additional timing cut is applied based on the phonon rise time. These cuts combined result in a misidentification probability of 10^{-6} for electron recoils. The neutron background was determined by Monte Carlo simulations of cosmogenic muons, spontaneous fission, and (α, n) ² processes. The CDMS II collaboration performed the above analysis blindly (that is, they did not look at any data in the expected WIMP region), and upon unblinding of their data run taken over four periods from July 2007 to September 2008, they find two possible WIMP events with recoil energies of 12.3 keV and 15.5 keV as shown in Fig. 1.6. There is a 23% probability that at least two background events pass all cuts. Thus the CDMS results cannot be interpreted as significant evidence for WIMP interactions.

1.5.3 Liquid noble gas detectors

Liquid noble gases are attractive media for dark matter detection due to their intrinsic scintillation properties, ease of operation, and low cost. Liquid noble gases, particu-

²The notation (a,b) indicates any nuclear interaction in which the incoming particle is a, and the outgoing particle is b.

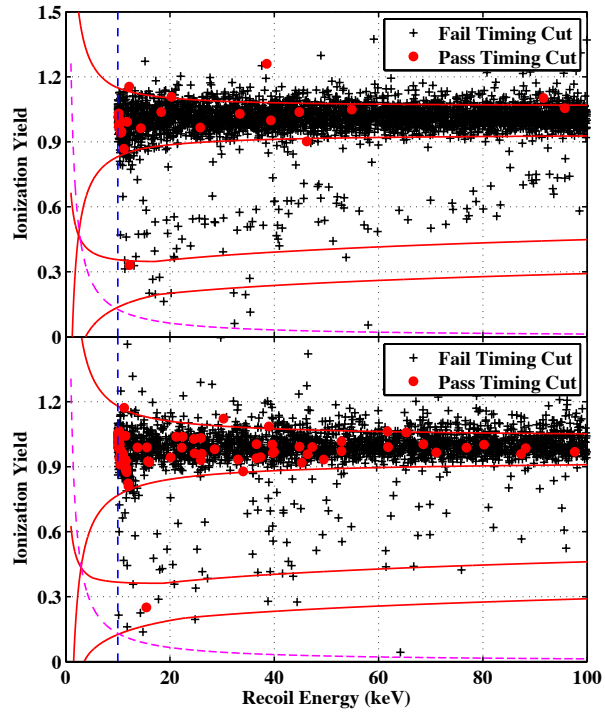


Figure 1.6: The results for two of the detectors in the CDMS II experiment. The upper bands are electronic recoils, while the lower bands are nuclear recoils. There are two events that pass the timing cut and are in the nuclear bands [30].

larly liquid argon, liquid xenon, and liquid neon, are unique in that they are very good scintillators, and they are transparent to their own scintillation light. This allows for scalability up to multi-ton masses which is necessary to achieve the sensitivity levels required for dark matter detection. Liquid scintillation provides background rejection by making use of the fact that electronic interactions and nuclear interactions produce differently shaped scintillation pulses. We will come back to “pulse shape discrimination” in studying the slow component lifetime. Pulse shape discrimination is practical only in liquid argon and neon; in xenon, the pulse shapes have lifetimes too close together to reliably distinguish them.

Liquid xenon and liquid argon are further attractive target media because they allow for ionization detection when an electric field is applied, as is done at the XENON (Sec. 1.5.3) experiment, among others. These detectors are usually dual-phase, with liquid mostly filling the detector volume and gas filling a thin layer at the top. The electric field drifts ionization electrons through the liquid, up to the gas phase, through which they are accelerated by a higher field. The accelerated electrons produce light by electroluminescence³, which is proportional to the amount of ionization. Thus ionizing radiation produces two signals: the first, denoted S1, is a scintillation pulse from the initial interaction with a target nucleus, and the second, denoted S2, is from the drifted electrons. Since the scintillation to ionization ratio in noble liquids depends on the energy loss per unit path length, dE/dx , of the incident particle [32], the S2/S1 ratio of each event provides very good discrimination between electronic interactions and nuclear interactions.

XENON

The XENON collaboration is the current leader in liquid noble gas detectors for direct detection experiments [33]. Located at Laboratori Nazionali del Gran Sasso in Italy, it is a three-dimensional position-sensitive, dual-phase (liquid/gas) cylindrical time projection chamber (TPC) filled with 161 kg of ultra-pure liquid xenon of which

³Electroluminescence is the process in which energetic electrons (in this case, the accelerated electrons), excite atoms by collisions, which produce light in a process similar to that of scintillation (see Sec. 2.2) [31].

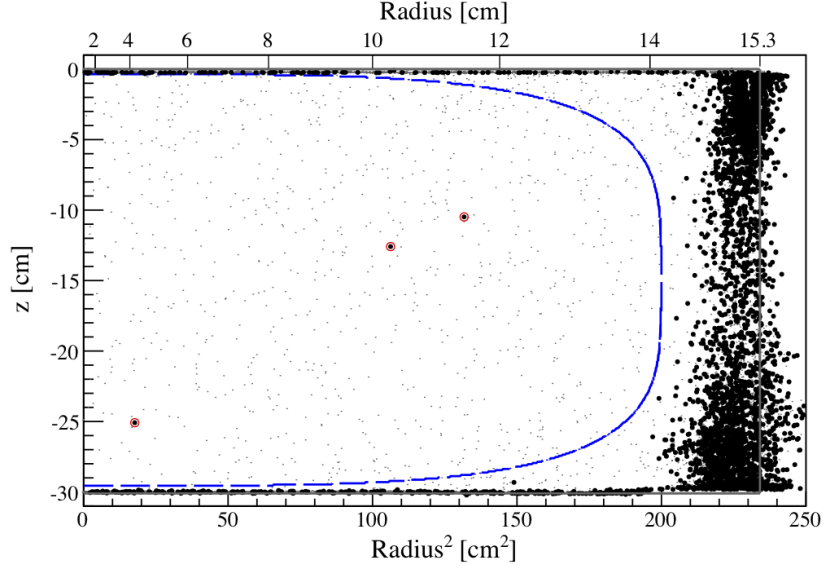


Figure 1.7: The dark matter run results from a 100.9 live-day data run of the XENON100 experiment. Three events pass all cuts with an expected background of (1.8 ± 0.6) events [34].

92 kg are used as an active veto and 62 kg are in the target volume. An array of 80 photomultiplier tubes (PMTs; see Sec. 2.4) at the bottom of the TPC is primarily for S1 detection while an array of 98 PMTs at the top is used for S2 detection and allows for (x, y) position reconstruction of an event with resolution < 3 mm. An electric field of about 530 V/cm parallel to the central axis of the detector drifts ionization electrons in the liquid xenon, extracts them to the gas phase, and accelerates them to produce S2. The timing between S1 and S2 provides < 2 mm resolution in z . The position reconstruction of events allows for the “fiducialization” of the detector in which only an inner volume of 40 kg is kept, where background events are drastically reduced. A run of 11.17 live days of background data, taken from October 20th to November 12th 2009, observed 22 events in the fiducial volume, but none of them fell within the acceptance region for the ratio of S2/S1. The XENON experiment recently completed a 100.9 live-day data run, taken between January and June 2010, finding after a blind analysis 3 events with an expected background of (1.8 ± 0.6) events [34]. The three events are shown in Fig. 1.7 with the fiducial volume denoted by the dashed blue line, and those events passing all cuts circled in red.

DEAP/CLEAN

The Dark matter Experiment with Argon and Pulse shape discrimination (DEAP) and Cryogenic Low Energy Astrophysics with Noble gases (CLEAN) family of detectors use only the scintillation channel of liquid noble gases for particle detection [35]. A series of successively larger single phase detectors is in development, which, without the use of drift field, will rely on pulse shape discrimination (see Sec. 2.2) for background rejection of electronic recoils. The first is DEAP-1, which has an active volume of 7 kg of liquid argon and is being used to validate single phase liquid noble gas detection technology in the search for dark matter [36]. The second is the MiniCLEAN experiment, a dual purpose detector that uses both liquid neon (85 kg fiducial volume) and liquid argon (100 kg fiducial volume) interchangeably to measure the pp solar neutrino flux as well as conduct a WIMP search [37]. The DEAP-3600 detector is under construction at the time of this writing and will have a 1000 kg fiducial volume of liquid argon [35].

Dark matter signature and DAMA/DAMA-LIBRA

Another direct detection technique is to detect the annual or diurnal (daily) modulations of the dark matter flux. There is an expected WIMP “wind” continuously hitting the Earth, and as the Earth rotates around the sun, the WIMP flux will vary sinusoidally, creating the so-called “annual modulation signature”, which reaches its maximum in June and its minimum in December [38, 39]. The DAMA experiment and its successor, DAMA/LIBRA, at the Gran Sasso National Laboratory are designed to detect this modulation, using radiopure sodium iodide doped with titanium, NaI(Tl), as the target scintillator [40, 41]. The detected signal must satisfy a stringent set of requirements to be considered the WIMP signature: (i) since the probability of a WIMP interacting multiple times in a detector is negligibly small, the modulation should be in the single hit rate, which is defined as events where only one interaction is detected; (ii) events should be in a well defined low energy region corresponding to the expected nuclear recoil energy for a WIMP; (iii) the signal should follow a cosine-

like behavior; (iv) the period should be one year, (v) with a phase that peaks in June, and (vi) with limited amplitude of $\leq 7\%$ in the region of maximal sensitivity. Very few effects can fulfill all these requirements; indeed, only systematic effects can mimic this modulation, and no other effects investigated thus far have so unambiguous a signature.

The DAMA/LIBRA experiment currently has the largest experimental exposure with 0.82 ton yrs. The results are shown in Fig. 1.8. However, no other experiment has been able to confirm DAMA's result [16], while other experiments seem to exclude the WIMP cross section that DAMA claims [13].

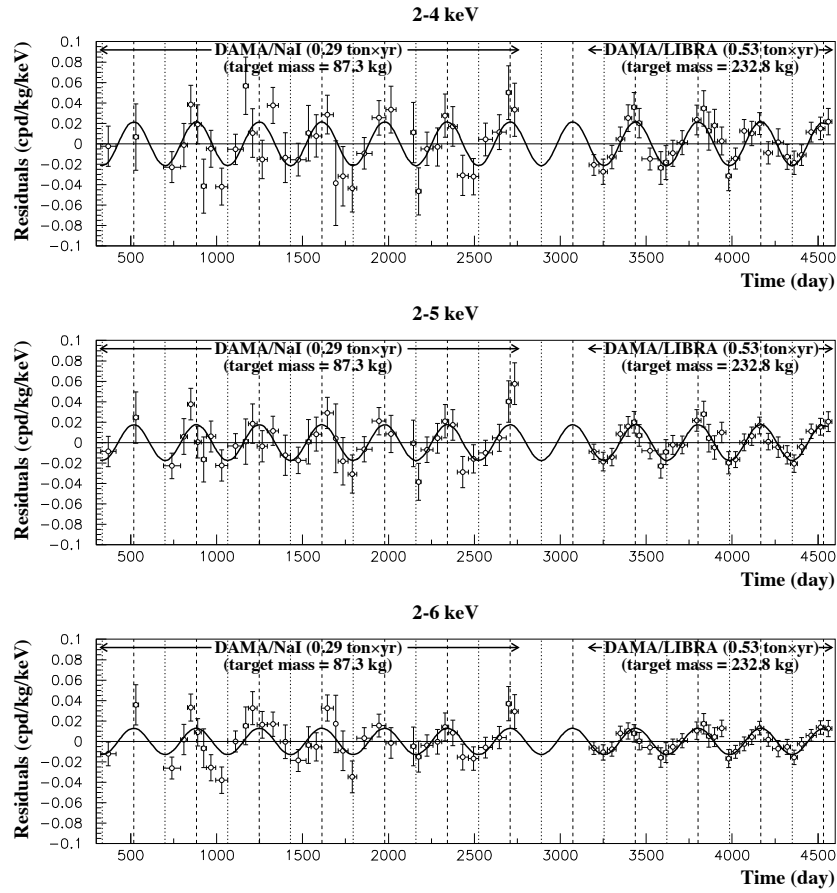


Figure 1.8: The results from the DAMA/NaI and DAMA/LIBRA experiments in various energy ranges.

Chapter 2

LAr TPC Experiments and DarkSide

The DarkSide experiment is a liquid argon scintillation and ionization detector being developed at Princeton University. At the time of writing, it is still in prototype and has just completed the second iteration of a 10 kg detector. The main feature of the second iteration is the installation of the electric field cage used to produce S2. DarkSide is developing a 50 kg detector for a first physics run before moving on to a multi-ton scale detector. We give here some of the main features of liquid noble gas detectors with a focus on liquid argon and its use in the DarkSide experiment.

2.1 Liquid argon

As mentioned in Sec. [1.5.2](#), one of the main techniques in direct detection experiments is the use of liquid noble gases as the target material for WIMP interactions. Liquid noble gases are well suited for direct detection experiments because they are transparent to their own scintillation light, allowing for the simultaneous detection of scintillation and ionization signals. An incident energetic particle, such as an energetic photon, neutron, or WIMP, produces scintillation light by one of the following processes. In the case of a photon, the incident gamma ray Compton scatters off an atom, releasing an energetic electron that is quickly stopped by nearby atoms. The

excited atoms release their energy as the S1 scintillation light (see Sec. 2.2). In the case of energetic neutrons and WIMPs, the incident particle scatters off a nucleus, imparting some of its kinetic energy to the atom. The energetic atom is also quickly stopped by nearby atoms, producing scintillation light. The energy deposition also ionizes many electrons that are drifted towards one end of the detector by a strong applied electric field, where they produce the S2 signal.

Liquid xenon and liquid argon are the most popular noble gases for dark matter detectors: liquid xenon is attractive for its low natural contamination levels, high density, and high ionization yield, while liquid argon is attractive for its lower cost and potential for pulse shape discrimination. Liquid argon differs from liquid xenon on three main accounts. First, while the 175 nm scintillation photons from liquid xenon are detectable by photomultiplier tubes, the 128 nm scintillation photons from liquid argon require the use of a wavelength shifter, usually tetraphenyl butadine (TPB), to shift the photons into the visible spectrum (wavelength 440 nm for TPB). Second, the characteristic decay times of the excited states of the atom are much more separated for liquid argon than for liquid xenon, allowing for pulse shape discrimination (PSD) of background events (see Sec. 2.2). Third, atmospheric argon contains natural ^{39}Ar contamination, which beta decays to ^{39}K with half life 269 yr, presenting a significant background for large scale detectors.

2.1.1 Depleted argon

One of the unique features of the DarkSide experiment is that it will use depleted argon rather than atmospheric argon in its active volume. Most other liquid argon based direct detection experiments use atmospheric argon, which contains trace amounts of the radioactive isotope ^{39}Ar , produced by cosmic ray interactions in the upper atmosphere [42]. The specific activity of ^{39}Ar is measured to be ~ 1 Bq/kg of atmospheric Ar [43]. At the multi-ton scale, this radioactivity will severely hamper the performance of the detector, and in this regard, liquid xenon is a favorable alternative for scintillation and ionization detectors. However, underground natural sources of ^{39}Ar -depleted argon have been discovered in the US National Helium Re-

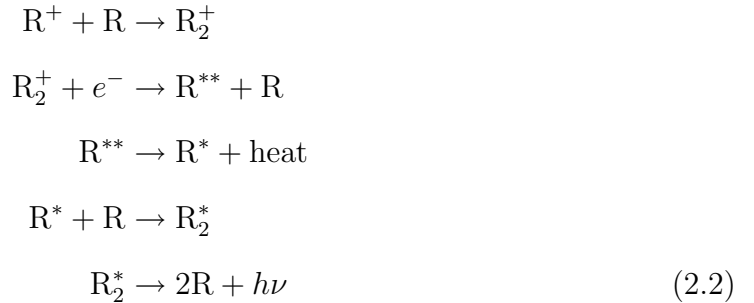
serve at the Cliffside Storage Facility outside Amarillo, TX [44]. This depleted argon is found to contain less than 5% the amount of ^{39}Ar compared to atmospheric argon. Although depleted argon was not used in the 10 kg prototype, it will be one of the main features in future versions of the DarkSide detector.

2.2 Scintillation detection

Scintillation is the luminescence emitted from a material in the presence of energetic particles. Scintillation in liquid noble gases, such as liquid argon and liquid xenon, arises from two distinct processes: the de-excitation of an excited nucleus, which in this section is generically denoted R^* , and the recombination of an ionized nucleus, R^+ , with an electron [45]. The processes are, respectively,



and



where $h\nu$ denotes the vacuum ultraviolet (VUV) photons emitted in the scintillation process with wavelength 128 nm. Now, the excited dimer, R_2^* , may exist either as a singlet or triplet state, and the de-excitation from each of these states back to the ground state occurs according to an exponential time distribution but on two different time scales with the singlet state decaying faster than the triplet state [46]. The difference between the two decay times varies for different liquid noble gases. For

liquid argon, the singlet state has a lifetime of $\tau_s \approx 7$ ns, while the triplet state decays slower with a lifetime of $\tau_t \approx 1.6$ μ s; for liquid xenon, the singlet and triplet lifetimes are $\tau_s \approx 4.3$ ns and $\tau_t \approx 22$ ns, respectively; and for liquid neon, they are $\tau_s < 18.2$ ns and $\tau_t \approx 14.9$ μ s [46]. For liquid argon and liquid neon, the many orders of magnitude difference between these two components, dubbed “fast” and “slow” for the singlet and triplet states, respectively, allows for reliable distinction between them. Note that this is not possible in LXe because the two decay components are too close to each other. The ratio of the intensities of singlet to triplet states depends on the value of dE/dx for the particular particle. Since dE/dx is different for recoiling particles produced by electronic processes as compared to nuclear processes, measurement of the ratio of the two allows for discrimination between different electronic and nuclear interactions [32]. A discrimination power, i.e. the ratio of mis-identified β events to the total, of $10^8 : 1$ for liquid argon has been demonstrated [47].

2.3 Ionization detection

Though pulse shape discrimination is a powerful tool for detecting different kinds of background, there is another effective technique at the disposal of some direct detection experiments which makes use of the ionization electrons produced by the energetic particle interactions. With the use of a large electric field applied across the liquid argon, the ionization electrons are drifted to one of end of the detector. They are then extracted to a gas phase by a higher electric field, where they produce a large amount of electroluminescence, creating the S2 signal. Due to their different ionization densities, the ratio of S2 to S1 is different for electronic and nuclear interactions, and so this parameter is also used as an effective way of discriminating between them. An additional discrimination power of $10^2 : 1$ has been achieved [47].

In addition to the ionization detection, the drift field allows for z positional reconstruction of an interaction event. Within larger detectors, x - y positional reconstruction is also possible. Depending on the strength of the field, electrons may take ~ 100 μ s to drift across the detector, whereas the S1 pulse has a maximum lifetime of

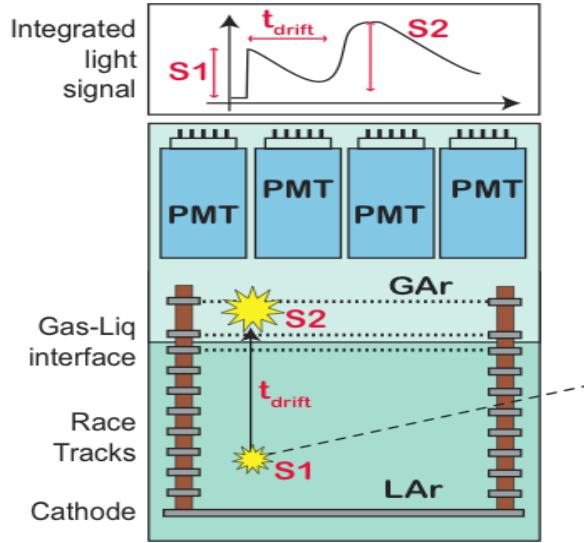


Figure 2.1: Schematic of LAr time projection chamber. [16] The race tracks in the figure refer to the electrodes used to produce the electric field across the detector.

a few microseconds. Taking the z -axis along the central vertical axis of the detector, the timing between S1 and S2 gives z -positional resolution, creating a time projection chamber. In addition, because the S2 signal is produced in a thin layer of gas directly below the top photomultiplier tubes (see Sec. 2.4), the electroluminescence is localized and the relative intensities of the top photomultiplier tubes allows for x - y positional reconstruction. Three-dimensional position reconstruction allows for fiducialization of the active volume, where, due to the self-shielding nature of the detector, the background is expected to be extremely low. That is, most particles are stopped near the edge of the detector volume, so the inner region is very quiet. Figure 2.1 shows a schematic drawn of a scintillation event in a LAr TPC. The “race tracks” refer to the electrodes used to produce the electric field across the detector.

2.4 Photomultiplier tubes

Photomultiplier tubes (PMTs) constitute the main means of detection of the scintillation and ionization signals. A photon that enters a PMT hits the photocathode, where, via the photoelectric effect, an electron is ejected with the energy of the incom-

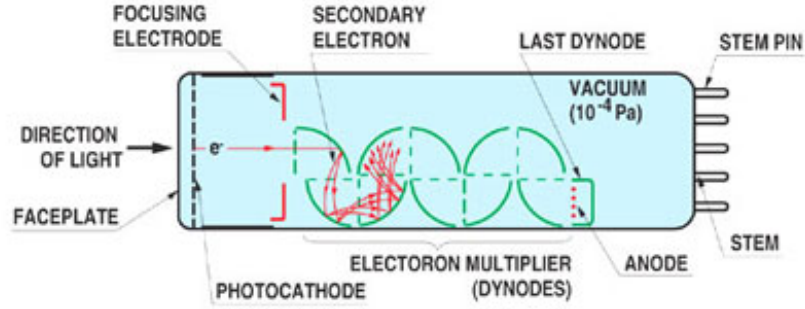


Figure 2.2: Schematic of a simple photomultiplier tube [49]

ing photon less the work function of the photocathode. These electrons are referred to as photoelectrons. The photoelectron is accelerated to the first dynode, where many electrons are released by the process of secondary emission¹. These electrons are accelerated to the second dynode, where more electrons are ejected and accelerated to the third dynode, and so on through a series of about a dozen dynodes. Thus a single photoelectron is amplified up to $\sim 10^6$ electrons, which is then converted to a current at the anode. See Fig. 2.2 for a schematic of a typical photomultiplier tube.

Due to the presence of U and Th impurities, standard photomultiplier tubes present one of the greatest sources of internal radiation for liquid noble gas detectors [14]. The DarkSide experiment will make use of state-of-the-art photomultiplier tubes, the Quartz Photon Intensifying Detector (QUPID), which have U/Th impurity levels a factor of ~ 100 lower than current low radioactive PMTs [50, 51]. Although these PMTs were not used in the second iteration of the 10 kg prototype, they will be used in the third iteration before moving on to the 50 kg version.

Figure 2.3 shows a diagram of the DarkSide 10 kg detector TPC. The active volume is 9.5 inches in diameter and 9.25 inches high and is viewed from above by seven 3" PMT and from below by a single 8" PMT [53]. The top PMTs are used for both S1 and S2 detection, while the bottom one is only for S1.

¹Secondary emission is very similar to the photoelectric effect, except that the photon is replaced by an electron. An energetic electron impacts a dynode surface, allowing a number of secondary electrons to escape [48].

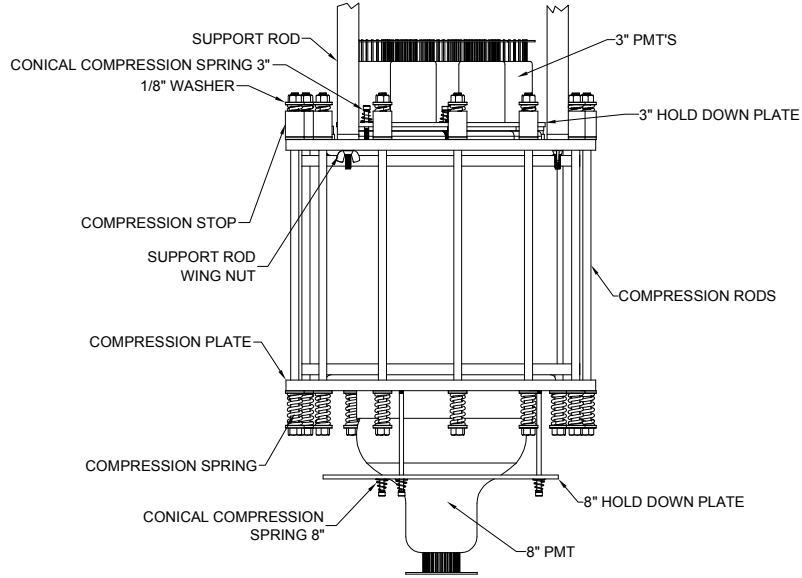


Figure 2.3: Schematic of the DarkSide 10 kg detector with various parts labeled [52]. There are seven 3" PMTs on top and one 8" PMT on the bottom viewing an active volume of 10 kg of liquid argon.

2.5 Active veto

The DarkSide experiment will make use of an active neutron veto consisting of liquid scintillator surrounding the active volume. The A series 100 PMTs will provide 6.5% coverage of the veto region. It will be used to tag neutron interactions by correlating in time scintillation events within the active region to a scintillation event caused by neutron scatter or capture on an atomic nucleus in the active veto region. Since the WIMP interaction cross-section is expected to be so small, the probability of it interacting more than once is negligible, so any simultaneous signal in both the active volume and the veto region will be rejected. Meanwhile, neutrons have a much higher probability of interacting more than once. Since neutron signals are indistinguishable from that expected from WIMPs, this will be one of the primary methods for identifying neutrons within the detector.

Chapter 3

Data and Cuts

The goal of this work was to get a general understanding of some of the first data from the second iteration of the 10 kg DarkSide detector with a focus on the slow-component lifetime for primary scintillation pulses. The main feature of this iteration was the installation of the electric field cage for the production of S2 signals. Thus most data from the detector has two prominent features: S1 and S2. In working towards an understanding of the slow component, we focus on S1, and to avoid dealing with complications arising from S2 signals appearing as S1, we restrict our focus to runs with zero electric field applied, referred to as null-field runs. Much of the following analysis arose from pursuit of the the slow-component lifetime, so we give first an overview of the scheme for its calculation, beginning with the physical data and lower level analysis before discussing the lifetime analysis in detail.

3.1 ^{22}Na coincidence runs

We measure the slow-component lifetime of electronic recoils using ^{22}Na coincidence runs. A ^{22}Na source is placed adjacent to the detector with an external counter, consisting of a sodium iodide (NaI) crystal scintillator viewed by a PMT, on the opposite side of the source from the detector. Sodium-22 beta decays to ^{22}Ne , producing a positron that quickly annihilates with an electron to produce two back-to-back 511 keV photons. These photons are emitted isotropically, so some of the photons will

align so that one of them hits the NaI counter outside the detector and the other photon will enter the detector, where it can Compton scatter off an Ar atom to produce an energetic recoiling electron. The electron does not travel far in the liquid argon as it quickly deposits its energy on many nearby Ar atoms either by excitation or ionization. These excited and ionized Ar atoms then scintillate as described in Sec. 2.2. The scintillation is then detected by a designated PMT, in this case the bottom 8" one, inside the detector. We trigger on the coincidence between the NaI counter and the 8" PMT at a rate of ~ 300 Hz. When triggered, we store the output of all the detector's PMTs from some time before the trigger to some time after. For null field runs, this window is typically $1 \mu\text{s}$ before the trigger to $7 \mu\text{s}$ after the trigger.

3.2 Data output and analysis modules

The data is sampled at every 4 ns, so that for a typical $8 \mu\text{s}$ time window, each event contains 2000 samples. Each sample is a discretized output from a digitizer, which divides the voltage output of the PMTs into 4096 bins. The digitizer has a range of 2V, so each bin represents ~ 0.5 mV. The digitizer also applies a large positive DC offset, and since each PMT gives a negative voltage output when a photon hits it, the corresponding sample is below the offset while still being positive. See for example Fig. 3.1, which shows a sample raw waveform. Thus for each event we have a discretized waveform both in time (typically in units of microseconds) and amplitude (typically in units of digitizer counts), and each PMT's output is represented by a drop in amplitude from a large positive offset. Figure 3.1 shows a typical S1 signal, where the fast and slow components are clearly distinguishable as the initial large spike and the proceeding smaller ones, respectively. Each smaller spike corresponds roughly to a single photon, but since each photon produces a single photoelectron in a PMT, we often work in terms of photoelectrons rather than photons.

For each coincidence run, we take data from the detector over some specified time. For a typical trigger rate of ~ 300 Hz for ^{22}Na coincidence runs, we record between 10,000 and 100,000 events per run. The digital waveforms for each event are then

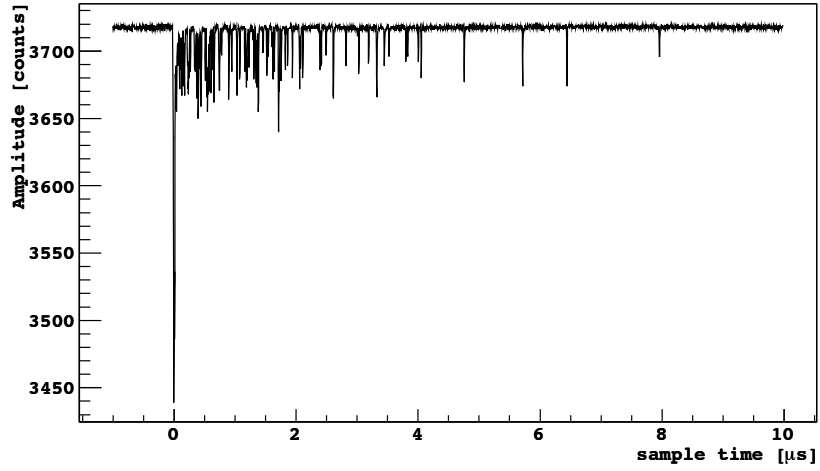


Figure 3.1: An example raw waveform for a scintillation pulse. The initial large spike is the fast component, while the rest of the spikes constitute the slow component.

analyzed by a series of modules, each dedicated to a different task. We give here a brief overview of the main modules used in the analysis for the calculation of the slow-component lifetime.

All analysis is done within the ROOT framework, a programming language based in C++ developed for particle physics. Its main advantage is to be able to handle large amounts of identically structured information, where each data point corresponds to a single event. All of the events are stored in a “Tree” and each of the modules then works on the Tree. The basic flow of the analysis is as follows: convert the data into a usable form, sum up the waveforms across all 8 channels, calculate the baseline on the summed waveform, and look for various analysis parameters of interest, including the locations of scintillation pulses (S1), their integrals, relative intensities, etc.

ConvertData

The first task is to convert the digitized data into a vector waveform and store alongside it various bits of useful information including the timestamp of the event and the channel from which the data came.

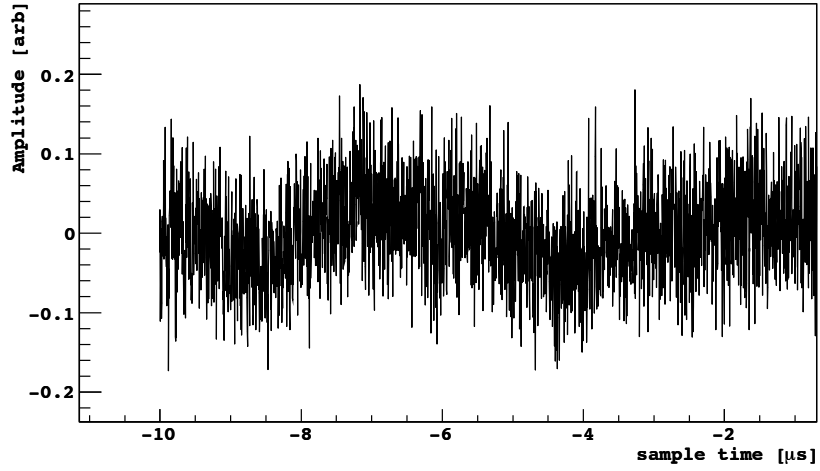


Figure 3.2: The sinusoidal fluctuation in the baseline of the summed waveform.

SumChannels

For the purposes of much of the downstream analysis, the waveform of each channel is scaled by the measured mean single photoelectron response (see Sec. 3.3) and added together. This summed waveform is then appended to the list of waveforms for the individual channels and is treated identically unless otherwise specified.

BaselineFinder

Each waveform has a large DC offset, called the baseline, which we wish to subtract before continuing with further analysis. For each event, we record between 1 and 10 μs worth of raw data before the trigger, known as the pre-trigger region, for the purpose of determining the baseline. In low electronic noise conditions, the baseline is constant throughout the waveform, and its value in the pre-trigger region determines it across the entire event. We would then determine the baseline by a simple averaging of the waveform values for some set of samples in the pre-trigger region. But the assumption of constancy is not valid: there are long timescale fluctuations in the baseline data, which follow a sinusoidal shape with period $\sim 5 \mu\text{s}$ as can be seen in Fig. A.1. See Appendix A for a more detailed characterization of this fluctuation. If not accounted for, these fluctuations will make downstream analysis difficult. Therefore, we use a

“moving baseline,” that uses only local values to estimate the baseline. For quiet regions, i.e. where there are no photoelectron pulses or other significant deviations from some mean value, the ten nearest samples (40 ns) are used to calculate a given sample’s baseline value. Any regions that are not quiet—that is, the RMS of the local region does not pass under a specified threshold—are assigned a value by interpolating between the two nearest samples for which a baseline was found. Thus the baseline is locally defined at all samples in the waveform, so it follows any fluctuations in the data that are of microsecond scale or longer. The entire waveform is then subtracted from this baseline, and most analysis proceeds from there.

Integrator

Given the baseline-subtracted waveform, we calculate its integral from start to finish by summing up all its values. This is denoted by the blue curve in Fig. 3.3. Since the waveform consists of a set of amplitudes (digitizer counts) over a set of samples in time, the integral is typically in units of count-samples. Given that there are ~ 0.5 mV per digitizer count and 4 ns per sample and that the PMT output passes through a $25\ \Omega$ resistor, we can convert the count-samples to Coulombs, and the integral then represents the amount of charge collected. For this analysis, we work in count-samples.

Note that the integral is quite steady in quiet regions of the waveform. This is a result of the moving baseline described above. Under Gaussian electronic noise, the integral is expected to run away as \sqrt{n} .¹ However, the moving baseline suppresses the runaway. This effect is useful for the next module but hides features of the waveform that may be of importance in later analysis.

PulseFinder

The next step is to identify scintillation pulses and record various values to characterize them. There are several different algorithms for finding the pulses; described

¹The integral is represented as a sum, and for noise, each term in the sum is Gaussian distributed with variance, say, σ^2 . The sum of n Gaussians is itself a Gaussian with variance $n\sigma^2$. Thus for n terms in the integral, the error is proportional to \sqrt{n} .

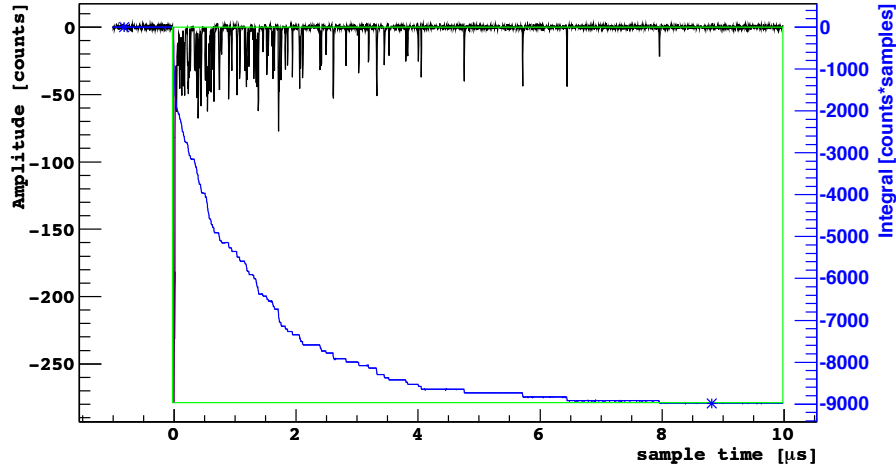


Figure 3.3: An example raw waveform with the baseline (red), integral (blue), and scintillation pulse (green) marked.

here is the “official” one. The main idea is to use the curvature of the integral as an indicator of the locations of each pulse. The integral curves much more steeply at the beginning of the pulse as the S1 signal appears. However, at a sample to sample scale (with 4 ns per sample), the curvature is a poor indicator of pulses because electronic noise fluctuations in the baseline can produce false positives. Thus we use a down-sampled version of the integral in which we take only every 250th sample of the waveform; this corresponds to one data point every microsecond. Then only large fluctuations in the waveform on the scale of a scintillation pulse will be noticeable in the down-sampled waveform. Once it is established that a scintillation pulse is in the waveform, the start of the pulse is identified by the sample just before the waveform crosses a specified threshold, and it is an easy task to find the rest of the parameters, including the peak amplitude, its integral, and the number of photoelectrons observed. The end of the pulse is found when the waveform flattens out (it has very small slope). An example of a baseline-subtracted waveform is shown in Fig. 3.3 with the baseline, integral (on a different scale), and scintillation pulse identified. Though the PulseFinder algorithm was originally tuned to work on the summed waveform across all the channels, it is easily retuned to find the scintillation pulses in single channels.

There are a total of eight channels, labelled 0 through 6 for the top seven 3" PMTs and 7 for the bottom 8" one, available for use in the analysis. However, towards the end of the second iteration, channel 6 experienced performance issues, producing significantly smaller signals than the other 3" PMTs. Thus it is excluded from the rest of the analysis discussed here. Channel 7 (8" PMT) is also excluded from the analysis since its light response is quite different from the other channels (3" PMTs).

3.3 Single photoelectron calibration

Another separate analysis on which the slow-component calculation relies is the single photoelectron calibration. For any set of coincidence runs, we periodically take a calibration run, in which we remove the ^{22}Na source while an LED emits photon pulses into the detector. The pulses are attenuated such that we can measure the PMT response to small numbers of photons (roughly one or two), which we use to calibrate the size, or more specifically the integral, of single and multiple photoelectron signals. An example spectrum of these integrals is shown in Fig. 3.4. The calibration analysis fits single and multiple photoelectron size distributions to this spectrum. In particular, the mean and width of the single photoelectron integral distribution, denoted μ_{pe} and σ_{pe} , respectively, may be determined. The single photoelectron is shown as the center pink curve in Fig. 3.4.

3.4 Slow-component lifetime

The slow component is exponentially distributed with mean $\tau_t \approx 1.6 \mu\text{s}$ for liquid argon. Thus most of the analysis focuses on the slow-component region, roughly between 0.2 and 8 μs , which is well after the fast component ($\tau_s = 7 \text{ ns}$) and includes multiple slow-component lifetimes. One method for estimating the slow-component lifetime is quite simple: find the average waveform of many events and fit an exponential function to the slow-component region. But the devil lies in the details. The first task is to subtract the baseline from the raw waveforms. Though we can find an

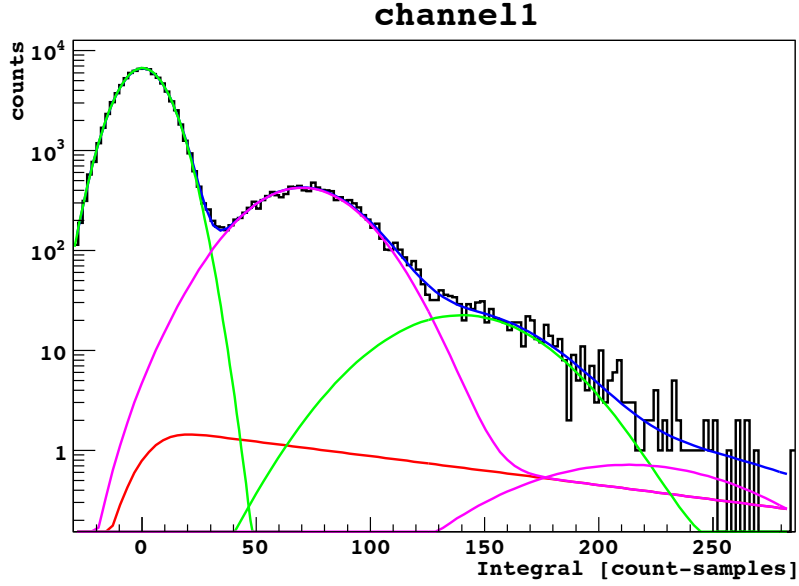


Figure 3.4: Example photoelectron pdf from a laser run. The single photoelectron pdf is indicated by the pink center pink line.

average waveform using the un-subtracted versions, many of the downstream tasks require that the baseline be subtracted. Furthermore, with a proper understanding of the baseline, the baseline-subtracted waveform should, in principle, be well fitted by a simple exponential without a constant background term. The next task is to exclude events that do not exhibit a typical slow component, including cosmic-ray muons, falsely triggered events, and events whose scintillation pulse starts at some time other than the trigger time. This chapter is concerned with such cuts as well as the baseline subtraction. We must then be able to attach errors to each average value in the waveform, which is discussed in detail in Chapter 4. Finally, we can fit an exponential function to the average waveform. We use a χ^2 minimization for the fitter. Thus it is important to have a detailed understanding of both the fitting function its errors. If the data is not completely exponential, we then pursue reasons why, as discussed in Chapter 5. The above program leads to discussions of afterpulsing, baseline calculation, and the behavior of null-field events outside the slow-component region, which are interspersed throughout.

3.5 Baselines

As described in Sec. 3.2, given a set of events in a run, the first step is to subtract the baseline to produce a “baseline-subtracted waveform;” we refer there for a description of how the moving baseline is calculated. We note here though that the moving baseline is not ideal for an error estimate in the average waveform calculation. The motivation of the moving baseline was to account for the long timescale fluctuations noticeable in quiet regions of the waveform. However, these fluctuations are also present beneath the signal regions of the waveform, so they should be included in the uncertainties attached to each point in the average. Using the moving baseline masks the sinusoidal fluctuations and whatever other features it takes out, so when adding the baseline uncertainties with the signal region uncertainties, the total uncertainty may be artificially low. Thus a “fixed baseline” calculation, where the baseline value is taken to be constant and is determined by a simple average of the waveform values in the pre-trigger region, may be necessary.

Indeed, the analysis done here uses a hybrid version of the two baseline calculation methods. We use a moving baseline to first calculate a baseline-subtracted waveform. This allows the `PulseFinder` module to find scintillation pulses effectively, which are necessary for various cuts on the events to be included in the average (see Sec. 3.6). Once all cuts are applied, we go back to the raw un-subtracted waveform of each surviving event to calculate a fixed baseline, which is then subtracted to obtain a new baseline-subtracted waveform. We find the average of these re-subtracted waveforms.

The fixed baseline calculation is not a blind averaging procedure. While the moving baseline can handle the presence of any pulses in the pre-trigger region by interpolating beneath them, a simple fixed baseline calculation will be insensitive to that and will be pulled away from the true baseline level. Thus we use a walking baseline method where we increment the length of the baseline by 15 samples at a time, checking that each increment is sufficiently quiet, i.e., that its RMS is below a specified threshold. The baseline is accepted only when sufficiently many samples are included (minimum of 50 samples) but extends until it finds a non-quiet region. The

fixed baseline value is calculated from the straight average of those samples. If no valid fixed baseline could be found, the event is excluded from the rest of the analysis.

3.6 Cuts

In order for the estimation of the slow-component lifetime to be believable, we must exclude events that are not primary scintillation pulses or that will adversely affect the average. One must always be wary in cutting events because doing so introduces some degree of bias, and it is easy to manipulate the data to say anything we want.

One of the most obvious cuts, then, is to exclude any events for which no baseline could be found. These are typically cosmic-ray muon events, which produce very large amounts of light within the detector. Since the analysis here is based on ^{22}Na coincidence runs, there are very few muon events, and the baseline cut generally excludes no more than one or two events. The second obvious cut is to exclude events whose waveform saturated the digitizer. These are also typically muon events. Next we cut events where `PulseFinder` failed to find any pulses. These are typically falsely triggered events, where no scintillation was actually present. Then we exclude events whose scintillation pulse was found, but which we do not want to include in the analysis, such as pulses that do not start at the trigger time. By design, the trigger time is very close to 0 for each event, so that the typical S1 pulse also starts at time 0. We also cut events that reach their peak amplitude too long after the start of the pulse. The rise time of the fast component is very fast –less than a few nanoseconds. Thus if the waveform reaches its peak amplitude significantly after the pulse’s start time, then it is not a typical S1. Furthermore, we expect the first pulse found the `PulseFinder` to be the S1 pulse. If this is not the case, then `PulseFinder` is not tuned properly. Thus we require that the first scintillation pulse have a start time after $-0.1\ \mu\text{s}$ and a peak time before $0.1\ \mu\text{s}$ relative to the trigger. As can be seen in Fig. 3.5, where the first pulses’s start times and peak times for all events in a run are histogrammed along the x - and y - axes, respectively, the majority of events pass this cut.

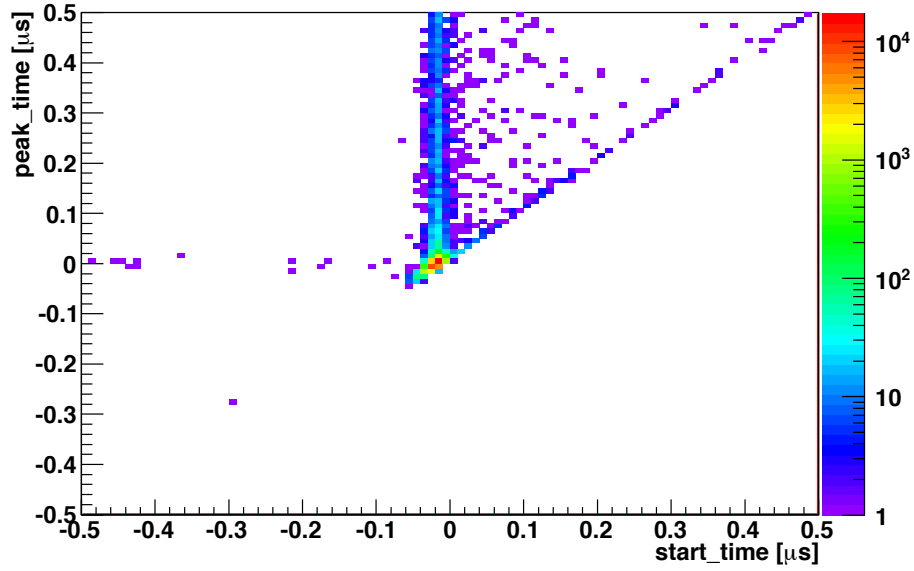


Figure 3.5: A histogram of first pulse peak times vs. start times. The bulk of events fall within a -0.1 to $0.1 \mu\text{s}$ window.

The most restrictive cut used in this analysis is the scintillation pulse peak amplitude cut. We exclude those events with peak amplitudes that are too small or too large. A sample spectrum of the first-pulse peak amplitudes is shown in Fig. 3.6. To determine which part of the spectrum to include, we look at how the slow-component lifetime behaves with different sections of the peak amplitude spectrum. (See Chapter 5 for a discussion of how we obtain the slow-component lifetime from a fitter.) We make 500-count wide cuts in peak amplitude starting with 0-500 counts, then 250-750 counts, and so on up to 1000-1500 counts, finding the slow-component lifetime for each cut. Note that the resulting lifetimes are not yet realistic because not all the parameters for fitting have been settled, but since we are looking only for a correlation in the behavior of the lifetimes with respect to cuts in pulse peak amplitude, we do not require the fitter to determine realistic lifetimes. The results for channels 0 through 5 are plotted in Fig. 3.7 where the x -axis shows the lower bound of each cut in peak amplitude and the y -axis shows the fitted lifetimes. The lifetimes do not vary over a very large range, and the only consistent behavior across all the channels is that the lifetimes from the cut starting at 250 counts is higher than that starting

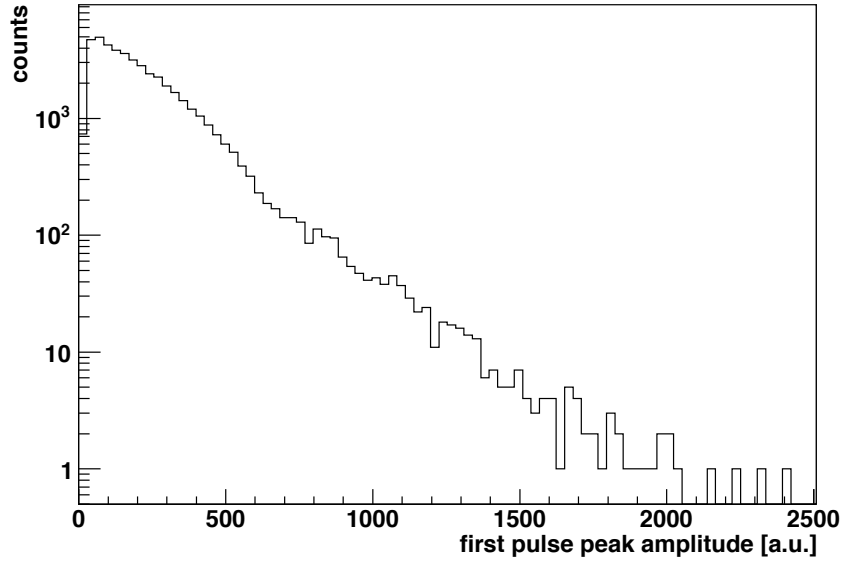


Figure 3.6: The spectrum of peak amplitudes of scintillation pulses in a ^{22}Na run.

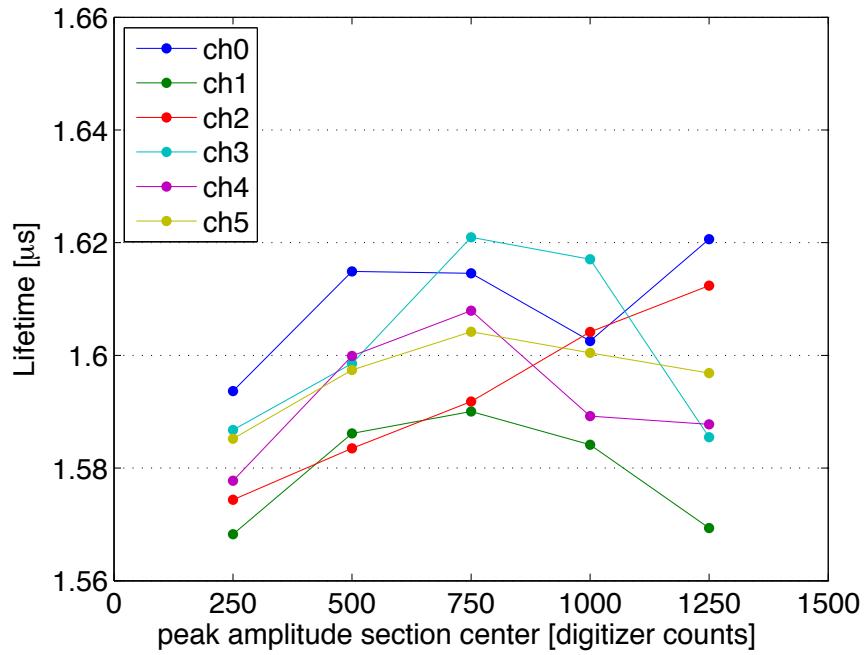


Figure 3.7: The fitted lifetimes for various sections in the pulse peak amplitude spectrum as shown in Fig. 3.6. The center value of each amplitude cut is plotted along the x -axis.

at 0 counts. Thus we take 250 to 1500 counts as our final pulse peak amplitude cut. The choice of upper bound is not as significant as the lower bound because there are few events with such large peak amplitudes.

Finally we apply a “spike” cut in which we exclude any events that have at least one spike, defined as any part of the waveform after the fast component that rises above a specified threshold value, typically $0.5 \times$ peak amplitude of the fast component. More specifically, we look after 200 ns from the start of the scintillation pulse, by which point the fast component is long over, for any part of the waveform that is more than half the peak amplitude of that component. This cut is to exclude events with afterpulses, which are caused by the ionization of residual gases in a PMT. The initial photoelectron may ionize a helium atom in the PMT as it is accelerated from the cathode to the first dynode. The ionized atom is then accelerated back to the cathode, where it produces many more photoelectrons, creating the afterpulse. The timescale for afterpulses ranges from several hundred nanoseconds to a few microseconds, exactly in the slow-component region, and their amplitudes are of the same scale as the fast component of a scintillation pulse [54]. Thus, if unaccounted for, afterpulses will artificially inflate the slow-component lifetime. By applying the spike cut, we find that afterpulses appear in 2% of events in all channels. Figure 3.8 shows the distribution of spike start times, which roughly agrees with that expected for afterpulses. Indeed, looking at individual events, these spikes visually match that expected for afterpulses: a single large amplitude pulse without any other structure associated. In particular there’s no slow component, so these aren’t additional scintillation pulses. Lastly, the effect of each of the cuts just described is summarized in Table 3.9.

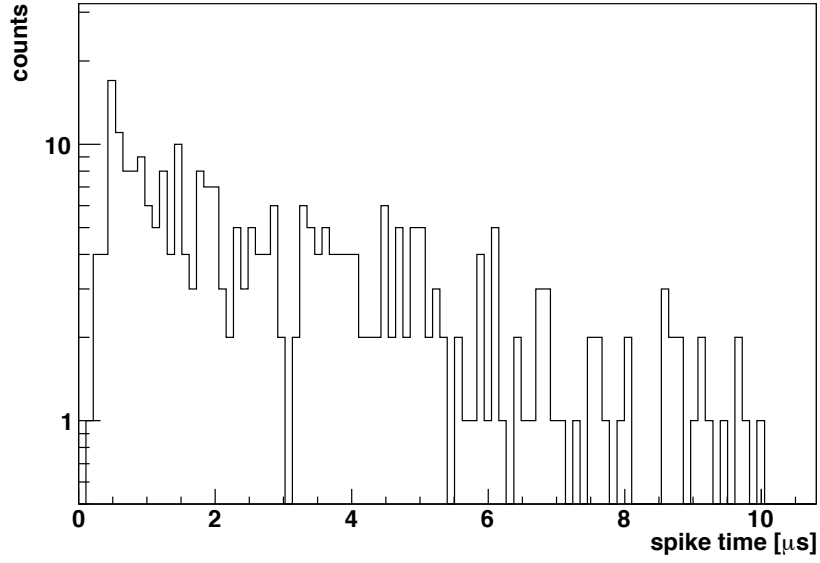


Figure 3.8: The distribution of spike times, which correspond to afterpulses.

Cut	Relative	Cumulative
Baseline	100%	100%
Saturation	99.85%	99.85%
Pulses>0	99.93%	99.77%
Timing	89.08%	88.88%
Peak amplitude	38.65%	34.35%
Spike	97.92%	33.64%

Figure 3.9: The relative and cumulative survival rates of each cut for channel 1 of run 1456. All events had good baselines.

Chapter 4

Photoelectron Counting Statistics

Given the events that survive all cuts, the next step is to calculate the average slow-component lifetime. We use a photoelectron counting argument to motivate taking a summed waveform of all events and performing the fits on the sum. The slow component lifetime is exponentially distributed. This means that given a large number of “slow-component photoelectrons”—i.e., photoelectrons due to photons produced by de-excitation of argon atoms in the triplet state—their cumulative waveforms will build up an exponential distribution. In a given scintillation pulse, there are hundreds to thousands of photoelectrons in the slow-component region (see Fig. 3.1). This is not enough to clearly define an exponential distribution; the waveform is still dominated by single (or double or triple) photoelectron waveforms. Thus, it will be difficult to fit an exponential function to a single scintillation pulse. So we build up the distribution using the slow-component photoelectrons from all events (that survived the cuts as described in Sec. 3.6) by summing the waveforms of each event.

In fact, it is more common to find an average waveform [46]. Indeed, taking the sum is essentially equivalent to taking an average. However, taking the average presents some complexities for determining the errors to be used in the fits. We present here an uncertainty calculation that is motivated by the photoelectron counting interpretation. First, we discuss how to implement the summed waveform.

4.1 Summed waveform

Let us denote the j^{th} sample of the i^{th} event's waveform by $y_i(t_j)$ in units of digitizer counts, which corresponds to sample time t_j relative to the event's trigger time. For a region of interest between times t_{j_1} and t_{j_2} , we first calculate the integral of the waveform, which in practice is the sum

$$y_i = \sum_{j=j_1}^{j_2} y_i(t_j) \quad (4.1)$$

Then the total number of photoelectrons for the i^{th} event is

$$n_i^{\text{tot}} = \frac{y_i}{C_{pe}} \quad (4.2)$$

where C_{pe} is a conversion factor for the number of count-samples per photoelectron as given by the calibration analysis described in Sec. 3.3. Although slightly unphysical, we may then interpret the number of photoelectrons at a sample t_j in the i^{th} event to be

$$n_i(t_j) = \frac{y_i(t_j)}{y_i} n_i^{\text{tot}}. \quad (4.3)$$

That is, the ratio of the waveform at sample j to the total integral is equal to the ratio of the number of photoelectrons at sample j to the total number of photoelectrons. This is unphysical because a single photoelectron is seen as a pulse over multiple (usually about 6-10) samples. Therefore it makes little sense to count the number of photoelectrons to arrive in a single sample. Nevertheless, this interpretation is justified by the fact that we will be summing over many events, where single photoelectron spikes get buried. This sum at the j^{th} sample, the “summed sample,” is

$$N(t_j) = \sum_{i=1}^n n_i(t_j) \quad (4.4)$$

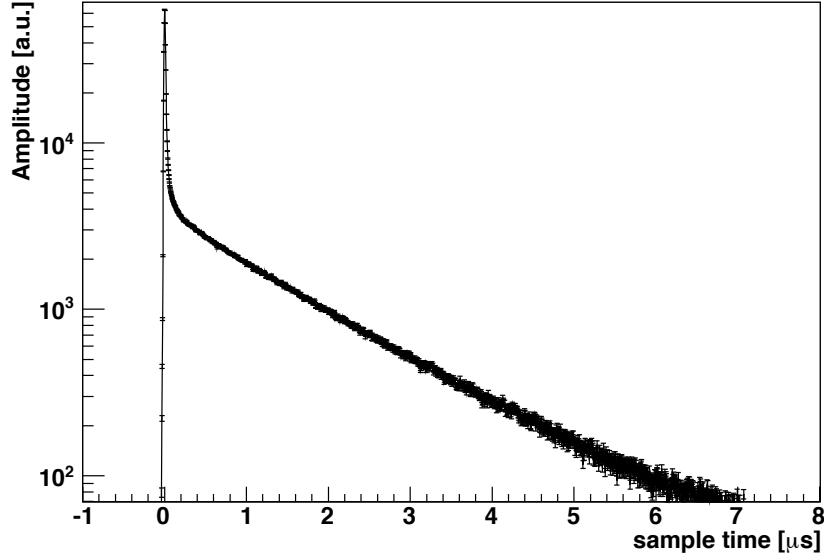


Figure 4.1: Summed waveform.

where n is the total number of events to be included in the sum, i.e., the number of events passing all cuts. Putting Eqs. 4.1, 4.2, and 4.3 into Eq. 4.4, we get

$$\begin{aligned}
 N(t_j) &= \sum_{i=1}^n \frac{y_i(t_j)}{y_i} n_i^{tot} \\
 &= \sum_{i=1}^n \frac{y_i(t_j)}{C_{pe}}
 \end{aligned} \tag{4.5}$$

Note that the summed sample, $N(t_j)$, is independent of the integral, y_i , for each event. Figure 4.1 shows an example summed waveform. The “summed waveform” here is not to be confused with the “summed waveform” of Sec. 3.2. The rest of this analysis is concerned with the summed waveform as defined by Eq. 4.5. We then fit this (huge) waveform, for which we must know the errors for each point.

We note here that another reason channel 6 is excluded from this analysis is that it has an unexplained structure at 2μ s in its summed (and average) waveform, as can be seen in Fig. 4.2. Before proceeding with the uncertainty calculation, we discuss another anomalous behavior of the summed waveform.

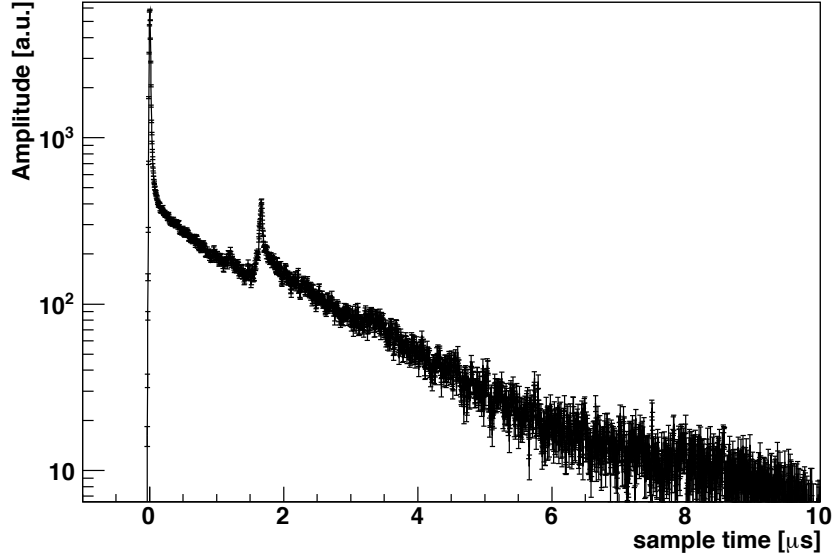


Figure 4.2: Channel 6 had significant performance issues, including overall smaller light response than the other channels and an unexplained structure at $2 \mu\text{s}$.

4.2 Alternating behavior

The baseline of the summed waveform exhibits an alternating behavior from sample to sample as can be seen in Fig. 4.3. We did not find any explanation for this behavior. Although the amplitudes of the alternations are small relative to the amplitude of the summed pulse, they are generally larger than the uncertainties associated with each point and are correlated with neighboring samples. This will adversely affect the fits, though we did not quantify this effect. By eye, the behavior has a period of four samples. We eliminate the effect of these alternations by splitting the waveform into eight pieces as described in Sec. 4.3. We need only be careful that we split the waveform into a multiple of four number of pieces, each of which will then have a baseline without this behavior.

4.3 Uncertainties

Given the summed waveform, the next step is to determine the uncertainty for each point. We want the uncertainty for each summed sample to reflect not only the count-

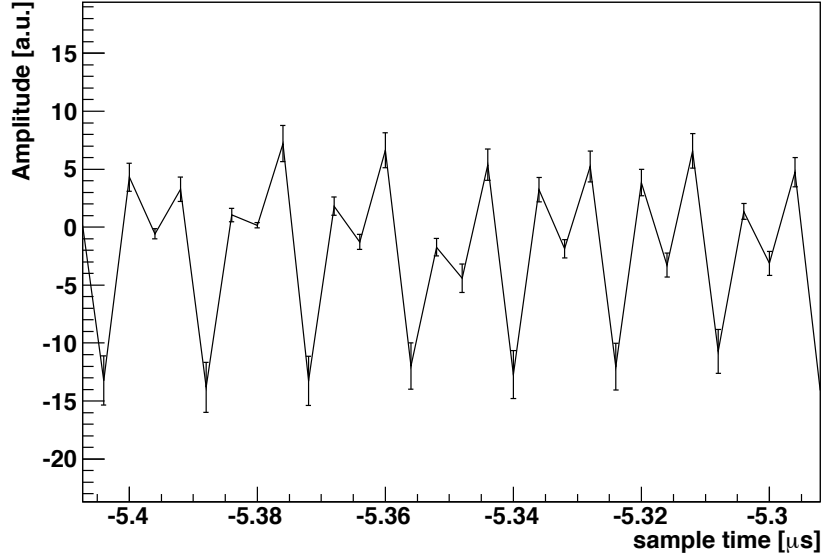


Figure 4.3: Alternating behavior in the summed waveform. This behavior is not discernible in individual events.

ing statistics of the photoelectron contribution to the sample but also the fluctuations due to both the varying size and nonzero pulse width of single photoelectrons, the last of which introduces correlations between neighboring samples.

4.3.1 Derivation

We estimate the uncertainties by the following model, which is motivated by the photoelectron counting interpretation described at the beginning of the chapter. We model the summed waveform as a very large set of photoelectrons arriving at some time t within a designated time window. Figure 4.4 shows a typical single photoelectron. They are quite uniform in shape but vary a lot in pulse height. Each photoelectron arrives with a size drawn from a single photoelectron probability distribution function with mean μ_{pe} and width σ_{pe} , to be determined in a separate analysis (see Sec. 3.3), and has some finite spread in time that affects neighboring samples of t . The spread is determined by the shape of the single photoelectron, which is determined by picking a typical single photoelectron at random. Since each photoelectron has some spread across multiple samples, there is some correlation between contributions

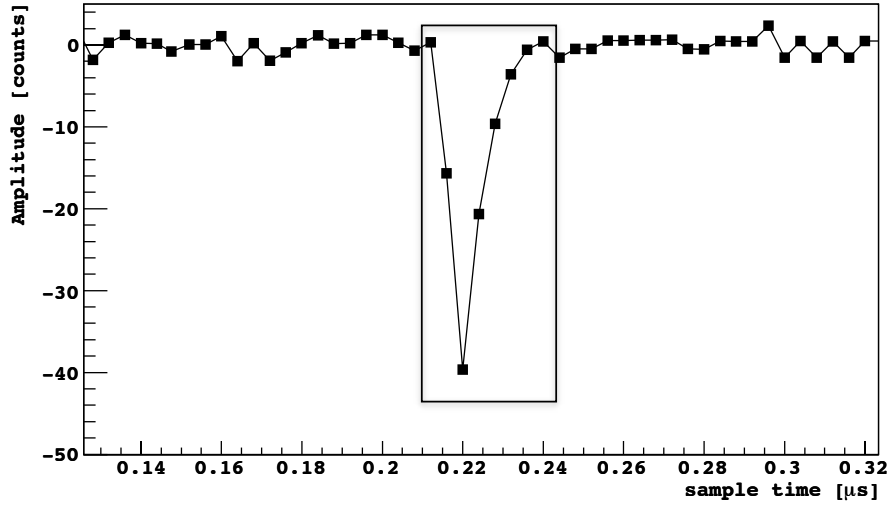


Figure 4.4: Single photoelectron shape with a bounding box around the 8 samples to characterize the photoelectron spread.

to adjacent samples in the summed waveform. We break this correlation by downsampling the summed waveform by a factor, l , chosen to be larger than the length (in samples) of the photoelectron pulse. We downsample the complete waveform l times, each with a different subset, or “split,” of the total. That is, each split takes every l^{th} sample with the first split starting at the first sample, the second split starting at the second sample, and so on up to the l^{th} split starting on the l^{th} sample. Thus in each split, neighboring samples are uncorrelated since a single photoelectron that arrives at sample j should have no effect on sample $j + l$. For a typical single photoelectron as shown in Fig. 4.4, the length is $l = 8$ samples (boxed off). We later fit exponential functions to all l splits.

We can calculate the uncertainties for the complete waveform (before downsampling) using a probabilistic argument as follows. We are interested in the errors of a single point within the summed waveform. We wish to be able to treat all l splits equally, so we note that since $l = 8$ samples spans 32 ns, on the scale of the expected lifetime at about $1.5 \mu\text{s}$, we can regard any single point to be residing in a locally flat region (within l samples away). This is the starting point for the model: we begin with a window of m samples, and photoelectrons arrive at any time within this

window with a uniform distribution. Let the random variable T_n = the arrival time of the n^{th} photoelectron. To begin, we model the photoelectron as occupying only a single sample time and account for its pulse width later. Then the probability that the n^{th} photoelectron arrives at the j^{th} sample time is

$$P[T_n = j] = \frac{1}{m}. \quad (4.6)$$

Now let $N_n(t_j)$ = the number of arrivals at sample j after n total photoelectron arrivals (across the entire window). Then the probability that there are k arrivals at sample j after n total arrivals is

$$\begin{aligned} P[N_n(t_j) = k] &= \binom{n}{k} \left(\frac{1}{m}\right)^{n-k} \left(1 - \frac{1}{m}\right)^k \\ &= \binom{n}{k} \left(\frac{n/m}{n}\right)^{n-k} \left(1 - \frac{n/m}{n}\right)^k \end{aligned} \quad (4.7)$$

In the limit as $n \rightarrow \infty$ with n/m fixed, Eq. 4.7 approaches the Poisson distribution and

$$P[N_n(t_j) = k] = \frac{e^{-n/m} (n/m)^k}{k!} \quad (4.8)$$

with expectation value and variance given, respectively, by

$$E[N_n(t_j)] = n/m \quad (4.9)$$

$$\text{Var}[N_n(t_j)] = n/m \quad (4.10)$$

Thus, in the large n limit, the number of photoelectrons at sample j is Poisson distributed with parameter $\lambda = n/m$.

Now each photoelectron will have a pulse height drawn from some single photoelectron distribution with mean μ_{pe} and width σ_{pe} as described in Sec. 3.3. Let $Y_i =$

the pulse height of the i^{th} photoelectron. Then

$$E[Y_i] = \mu_{pe} \quad (4.11)$$

$$\text{Var}[Y_i] = \sigma_{pe}^2 \quad (4.12)$$

for all i . We assume that Y_i is independent of Y_j for $i \neq j$, i.e., that the size of each photoelectron is independent of another; and that each Y_i is independent of $N_k(n)$, i.e., that the size of each photoelectron is independent of the number of photoelectrons to arrive at any sample. Then, still not accounting for the shape of the photoelectron pulse (that is, ignoring the spread of the photoelectron over l samples), the total pulse height at sample j after n total photoelectron arrivals is

$$Y = \sum_{i=1}^{N_n(t_j)} Y_i \quad (4.13)$$

Then using standard formulas for conditional expectations and variances [] we have

$$\begin{aligned} E[Y] &= E[E[Y|N_n(t_j)]] \\ &= E \left[E \left[\sum_{i=1}^{N_n(t_j)} Y_i | N_n(t_j) \right] \right] \\ &= E \left[\sum_{i=1}^{N_n(t_j)} E[Y_i | N_n(t_j)] \right] \\ &= E \left[\sum_{i=1}^{N_n(t_j)} \mu_{pe} \right] \\ &= E[\mu_{pe} N_n(t_j)] \\ &= \mu_{pe} \lambda \end{aligned} \quad (4.14)$$

and

$$\begin{aligned}
\text{Var}[Y] &= \text{Var}[\text{E}[Y|N_n(t_j)]] + \text{E}[\text{Var}[Y | N_n(t_j)]] \\
&= \text{Var}\left[\text{E}\left[\sum_{i=1}^{N_n(t_j)} Y_i | N_n(t_j)\right]\right] + \text{E}\left[\text{Var}\left[\sum_{i=1}^{N_n(t_j)} Y_i | N_n(t_j)\right]\right] \\
&= \text{Var}\left[\sum_{i=1}^{N_n(t_j)} \text{E}[Y_i | N_n(t_j)]\right] + \text{E}\left[\sum_{i=1}^{N_n(t_j)} \text{Var}[Y_i | N_n(t_j)]\right] \\
&= \text{Var}\left[\sum_{i=1}^{N_n(t_j)} \mu_{pe}\right] + \text{E}\left[\sum_{i=1}^{N_n(t_j)} \sigma_{pe}^2\right] \\
&= \text{Var}[\mu_{pe}N_n(t_j)] + \text{E}[\sigma_{pe}^2 N_n(t_j)] \\
&= \mu_{pe}^2 \text{Var}[N_n(t_j)] + \sigma_{pe}^2 \text{E}[N_n(t_j)] \\
&= \lambda (\mu_{pe}^2 + \sigma_{pe}^2)
\end{aligned} \tag{4.15}$$

Now we consider that each photoelectron has a finite spread. We assume that each photoelectron has the same shape as shown in Fig. 4.4. We denote the components of the single photoelectron shape by $\{c_k\}$, where $k = 1, \dots, l$ and the components are normalized so that $\sum c_k = 1$. Thus the total contribution of the photoelectron to all (neighboring) samples is given by the size drawn from the single photoelectron pdf. A single photoelectron with leading edge at time t_j will affect the samples not only at j but also $j+1, j+2, \dots, j+l$, with a contribution of $c_k Y_i$ at the $(j+k)^{\text{th}}$ sample.

Now, there is an equal probability that the i^{th} photoelectron will contribute $c_1 Y_i$, $c_2 Y_i, \dots$, or $c_k Y_i$ to the j^{th} sample. Or to put it another way, we may regard photoelectrons as arriving as one of l different types, N_1, \dots, N_l , and each photoelectron may be any one of these types with equal probability. Furthermore, photoelectrons of each type will have an energy drawn from its corresponding energy distribution, $Y_i^{(1)}, \dots, Y_i^{(l)}$, where $Y_i^{(k)} = c_k Y_i$. Each N_k is still Poisson distributed with mean n/m and independent of each other. Then the total energy at sample j after n total arrivals is

$$Y = \sum_{i=1}^{N_1} Y_i^{(1)} + \sum_{i=1}^{N_2} Y_i^{(2)} + \dots + \sum_{i=1}^{N_l} Y_i^{(l)} = \sum_{k=1}^l \sum_{i=1}^{N_k} Y_i^{(k)} \tag{4.16}$$

Then

$$\begin{aligned}
 E[Y] &= E \left[\sum_{k=1}^l \sum_{i=1}^{N_k} Y_i^{(k)} \right] \\
 &= \sum_{k=1}^l E \left[\sum_{i=1}^{N_k} c_k Y_i \right] \\
 &= \sum_{k=1}^l c_k \mu_{pe} \lambda \\
 &= \mu_{pe} \lambda
 \end{aligned} \tag{4.17}$$

and

$$\begin{aligned}
 \text{Var}[Y] &= \text{Var} \left[\sum_{k=1}^l \sum_{i=1}^{N_k} Y_i^{(k)} \right] \\
 &= \sum_{k=1}^l \text{Var} \left[\sum_{i=1}^{N_k} Y_i^{(k)} \right] \\
 &= \sum_{k=1}^l \text{Var} \left[c_k \sum_{i=1}^{N_k} Y_i \right] \\
 &= \sum_{k=1}^l c_k^2 \lambda (\sigma_{pe}^2 + \mu_{pe}^2)
 \end{aligned} \tag{4.18}$$

where in the second line of Eq. 4.18, we use the independence of the N_k and $Y_i^{(k)}$.

Thus, at a given sample j , the expected size of the summed waveform

$$s_j = \mu_{pe} \lambda_j, \tag{4.19}$$

where λ_j is the expected number of photoelectrons in that sample, and the expected variance at that sample is given by

$$\sigma_j^2 = \lambda_j (\sigma_{pe}^2 + \mu_{pe}^2) \sum_k c_k^2. \tag{4.20}$$

Thus s_j/μ_{pe} is estimated by y_j , and the uncertainty is then estimated by

$$\sigma_{y_j}^{pe} = \sqrt{\frac{y_j}{\mu_{pe}^2}(\sigma_{pe}^2 + \mu_{pe}^2) \sum_k c_k^2}. \quad (4.21)$$

Note that this is only valid with the knowledge that we will downsample the summed waveform.

Finally, for each sample in each split, we add the baseline RMS of the (downsampled) summed waveform in quadrature with $\sigma_{y_j}^{pe}$. Since we are using the re-subtracted waveforms in the sum as described in Sec. 3.5, all fluctuations are present in each event's baseline, including electronic noise and the microsecond timescale sinusoid, and their summed effect is incorporated into the baseline RMS. The sinusoidal fluctuations are washed out in the summed waveform since their phases are random, but they still contribute to the RMS. We assume these fluctuations, which we measure in the pre-trigger region, are also present in the slow-component and represent an independent source of error from the photoelectron counting statistics. Then the total uncertainty for each summed sample is

$$\sigma_{y_j} = \sqrt{\text{RMS}_{\text{base}}^2 + \frac{y_j}{\mu_{pe}^2}(\sigma_{pe}^2 + \mu_{pe}^2) \sum_k c_k^2} \quad (4.22)$$

where RMS_{base} is the baseline RMS for the corresponding split.

Chapter 5

Fit and Results

We fit the summed waveform to a simple exponential of the form

$$y = Ae^{-t/\tau} \tag{5.1}$$

with two free parameters to fit: amplitude, A , and lifetime, τ . The main tool for fitting a function to data is MINUIT, a numerical minimization program originally written in FORTRAN that has been repackaged into C++ and bundled into ROOT [55, 56]. It takes as input a set of data and a user defined function with a specified set of parameters. MINUIT then finds the set of parameters that minimizes the χ^2 statistic¹ between the data and the function. The main challenges here are to determine appropriate parameters for the fitter, including what the proper fit window should be and whether the inclusion of a background term is sensible or not. We also discuss several anomalies that reveal themselves in pursuit of these challenges.

¹The χ^2 statistic is an indicator of goodness of fit and is defined as

$$\chi^2 = \sum_i^n \left(\frac{y_i - f(x_i)}{\sigma_i} \right)^2$$

where y_i is the measured value at point x_i and has expected value $f(x_i)$ with standard deviation σ_i , and the sum is taken over all n points in the dataset.

5.1 Fit

We want to fit only the slow-component region, so we use a fit window of 0.5 to 6.0 μs as a rough first guess. Figure 5.1 shows such a fit, with lifetime $\tau = 1.564 \pm 0.004$ μs , and its residuals for a single split of a single channel. Also shown are the fitter’s numerical results, where neither the constant background term, “constant”, nor the “shot_noise” term were included in the fitted function. “Shot_noise” refers to an additional parameter in the fit function to account for noise. We find that such a term has little effect on the fitted lifetime, so we exclude it from the fits. The exponential does not fit well; notice that the χ^2 is large, and the data clearly deviates from exponential after ~ 6 μs . The next task then is to include the constant term, so the fit function is

$$y = Ae^{-t/\tau} + B. \quad (5.2)$$

The constant term is motivated by the nonzero summed waveform in regions well beyond the slow-component (see Sec. 5.4). The results are shown in Fig. 5.2, where the fitted lifetime is $\tau = 1.445 \pm 0.007$ μs . The fit is significantly better, though it deviates from the data late in the plotted region; we explore this behavior in Sec. 5.2.

We note that the RMS of the lifetimes across all splits for a given channel agrees well with the statistical uncertainties returned by MINUIT, as shown in Table 5.3. We note also that it is not very meaningful to calculate the average lifetime across all splits for a single channel since we expect them to be correlated, but since the RMS is small, we use the results for a single split as a reasonable representative of them all.

5.2 Anomalous structure

Looking out past 10 μs , we find not a simple deviation from exponential but an anomalous structure in the summed waveform spanning 0-60 μs with two peaks at ~ 20 μs and ~ 45 μs , as shown in Fig. 5.4. We have no tenable explanation for this structure, which appears consistently in all splits of all channels in many runs. Since

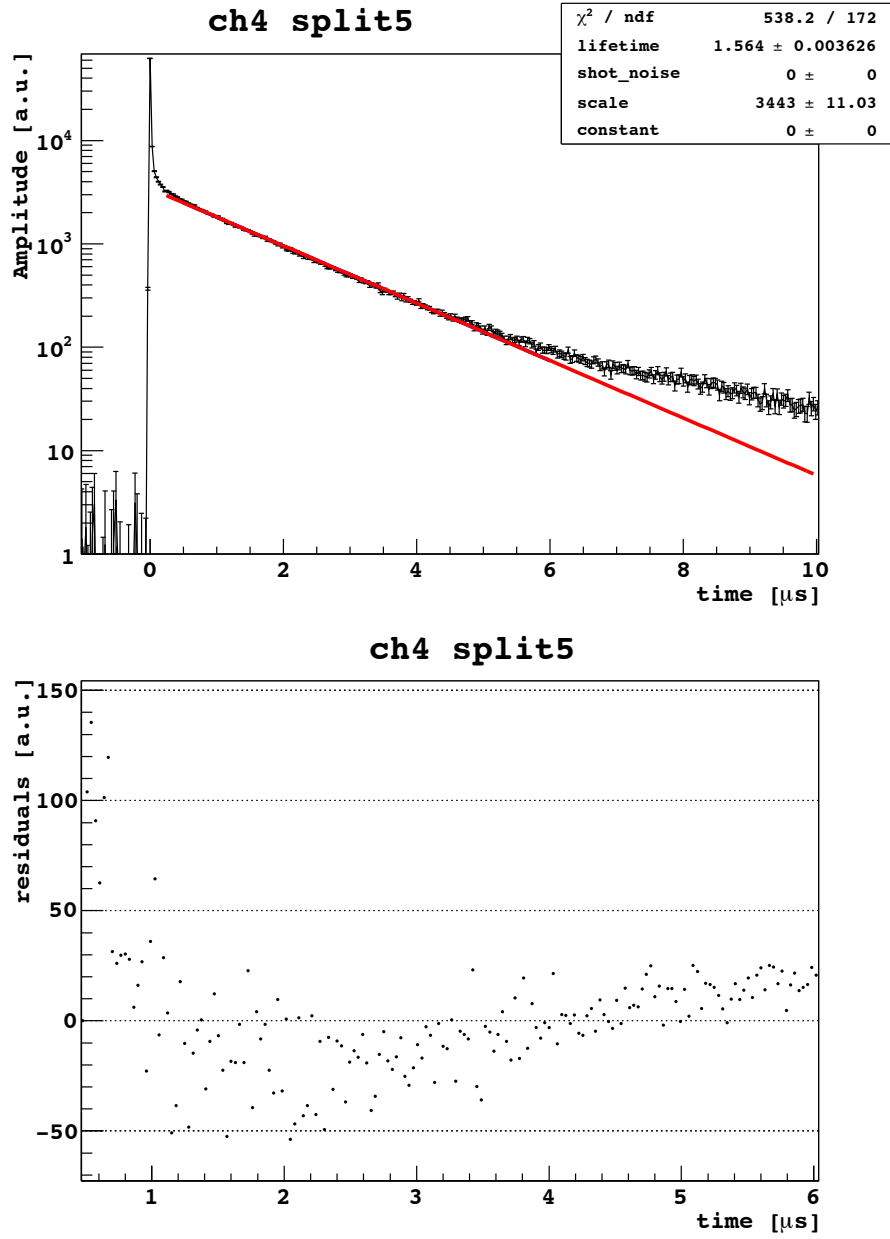


Figure 5.1: Sample fit (top) and its residuals (bottom) of the slow-component lifetime for a single split in a single channel with fit window 0.5 to 6.0 μs though the fitted function is plotted out to 10 μs . Shown also are the fitter results.

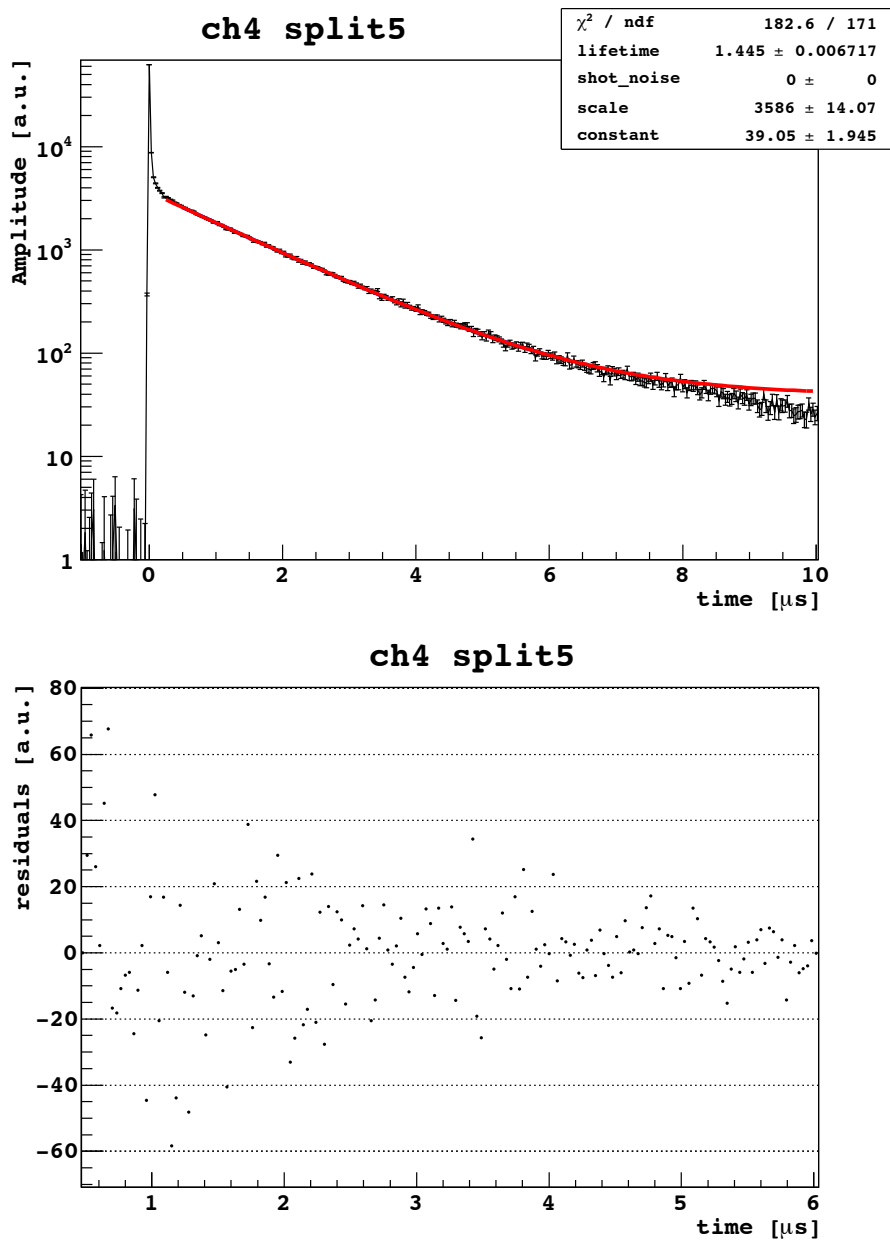


Figure 5.2: The same fit as Fig. 5.1 but with the constant term included. The fit window is still 0.5–6.0 μs , though the fitted function is drawn out to 10 μs .

Split	Fit
0	1.4798 ± 0.0060
1	1.4797 ± 0.0060
2	1.4735 ± 0.0058
3	1.4697 ± 0.0060
4	1.4636 ± 0.0058
5	1.4634 ± 0.0059
6	1.4667 ± 0.0058
7	1.4683 ± 0.0060
RMS	0.00612754

Figure 5.3: The fitted lifetimes across all splits for a single channel.

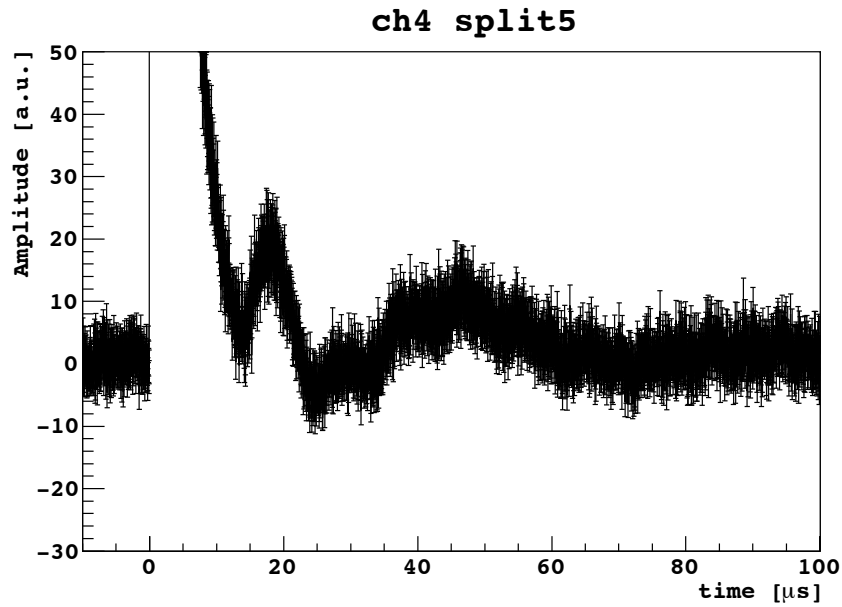


Figure 5.4: Unexplained structure appears in the summed waveform with peaks at $\sim 20 \mu\text{s}$ and $\sim 45 \mu\text{s}$.

the structure appears mostly beyond the slow-component region, it does not have a significant effect on our analysis. Indeed the deviation from exponential in Fig. 5.1 is likely the beginning of this structure. Thus we take only up to $6 \mu\text{s}$ as our upper limit of the fit window in order to avoid it.

5.3 Fit window

Motivated by the inability of an exponential fit without a constant term to account for all the data in a 0.5 to $6 \mu\text{s}$ window and as a check on the stability of the lifetime with the constant term on, we explore how the fitted lifetime behaves as a function of the fit window. We show here the fitted lifetimes for only a single split in each of the channels, leaving off error bars since they clutter the plots; but we keep in mind that the fit errors are typically $< 0.01 \mu\text{s}$. First, without the constant term, we fit the lifetime to various $2 \mu\text{s}$ sections of the slow-component region, as shown in Fig. 5.5. The fits for a single split in each of the six channels used in this analysis are plotted. The lifetimes show no sign of stability, and they increase steadily as we go further out in time starting from $\tau \approx 1.5 \mu\text{s}$ to $\tau \approx 2 \mu\text{s}$. This is further indication of the need for a constant term in the fit. Figure 5.6 shows the same fits with the constant term on. The lifetimes are significantly more stable: away from the upper edge of the slow-component region, the fitted lifetimes vary between $\tau \approx 1.4 \mu\text{s}$ and $\tau \approx 1.7 \mu\text{s}$. The reduced χ^2 are mostly close to unity, as shown in Fig. 5.7. We can further vary the fit window, changing the upper limit while keeping the lower limit fixed and vice versa, to see how the lifetime behaves; we find that the results are much in line with Figs. 5.5 and 5.6. Thus, while we believe MINUIT is finding good fits to the data, the variations in the fitted lifetimes are too large to give an absolute value to the slow-component lifetime.

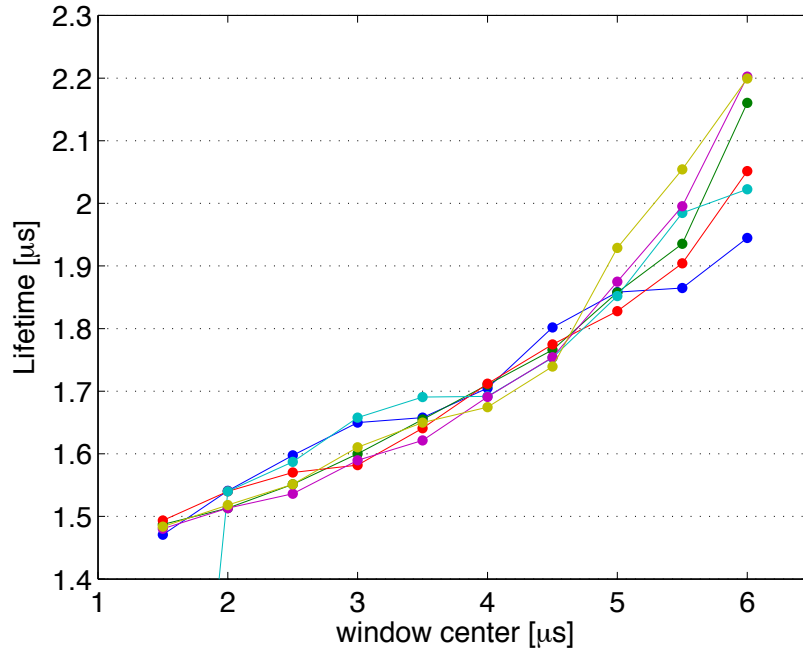


Figure 5.5: Fitted lifetimes, without the use of a constant term, to various sections of the slow-component region. Each section is $2 \mu\text{s}$ wide with the center of the each window plotted along the x-axis.

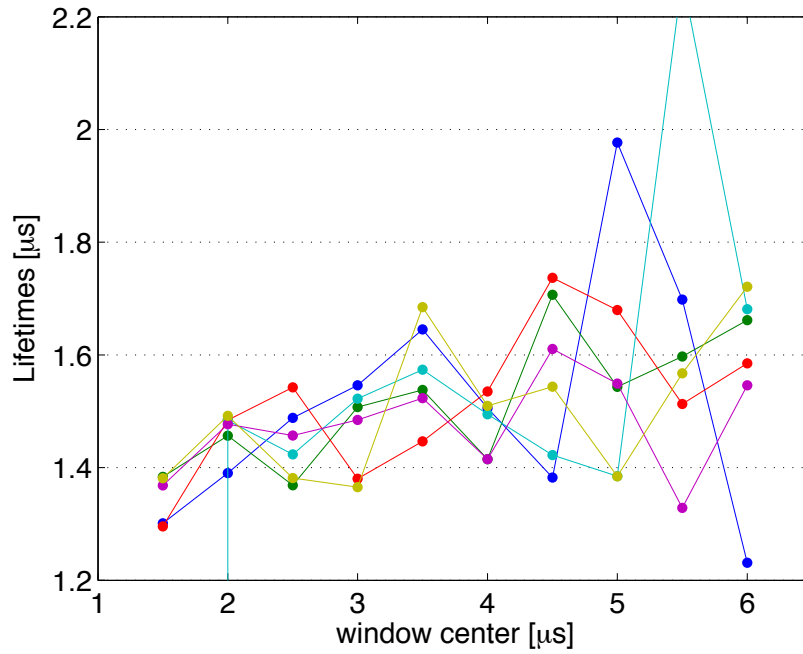


Figure 5.6: Fitted lifetimes, with the constant term on, to various sections of the slow-component region. Each section is $2 \mu\text{s}$ wide with the center of the each window plotted along the x-axis.

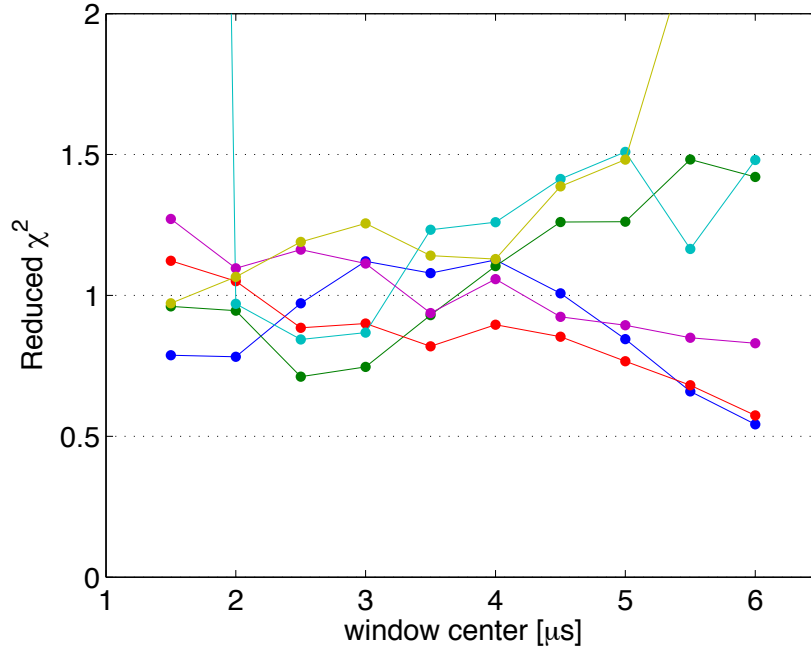


Figure 5.7: The reduced χ^2 values for the fits shown in Fig. 5.6.

5.4 Pre-trigger vs. post-pulse

In determining the summed waveform for our slow component lifetime calculation, we find that it does not return to zero amplitude after the scintillation pulse. This is not very obvious in looking only at the slow-component region ($< 8 \mu\text{s}$), but it is strikingly apparent when looking at the full average waveform of a long time window as shown in Fig. 5.8. The pre-trigger region is very close to zero amplitude, while the post-pulse region ($> 10 \mu\text{s}$ after the trigger) has a significant offset from zero and is much noisier. The main explanation for this behavior lies in the nature of the trigger mechanism for coincidence runs. As described in Sec. 3.1, we trigger on the coincidence of the 8" PMT and NaI discriminators, where we then collect data both before and after the trigger. Whenever the 8" PMT discriminator fires, it also produces a $60 \mu\text{s}$ veto on both the NaI and 8" PMT discriminators, regardless of whether the NaI discriminator fired. However, in general, the NaI discriminator fires before the 8" PMT, so by the time the veto starts, the NaI and 8" PMT discriminators will both have fired, and the coincidence gate will have triggered. Thus we still collect

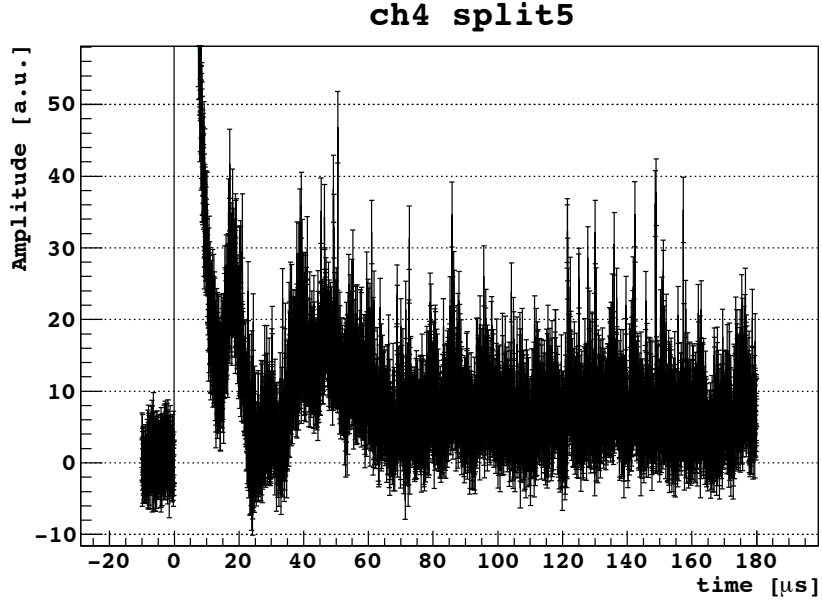


Figure 5.8: The summed waveform for a null field run with an unusually long data window. The pre-trigger region is close to zero and quiet, but it does not return to zero amplitude even long after the scintillation pulse dies out and is much noisier.

data as required, and any random pulse appearing in the detector alone, and not in the NaI, will induce a $60 \mu\text{s}$ veto. This mechanism ensures that the pre-trigger region is always quiet.

After the trigger, another ^{22}Na atom may decay and produce a scintillation pulse which will appear in all the detectors. When summing all the waveforms, these large pulses pull the baseline upward, accounting for the offset and noise in the post-pulse region of long time window runs. We can estimate the rate of post-trigger scintillation pulses by looking for spikes, as described in Sec. 3.6, out to the end of the waveform rather than just to the end of the initial scintillation pulse. We find the post-pulse spike times to be roughly uniformly distributed as shown in Fig. 5.9. with spike rates to be $\sim 750 \text{ Hz}$ for a spike finder threshold of $0.5 \times \text{peak amplitude}$ and $\sim 1750 \text{ Hz}$ for $0.3 \times \text{peak amplitude}$. The measured firing rate for the 8" PMT is higher at 1-3 kHz for ^{22}Na runs. We can reconcile the discrepancy between the spike rate and the measured activity by noting that the spike finder is applied after all previous cuts, particularly the peak amplitude cut. Thus the spike finder is biased towards larger

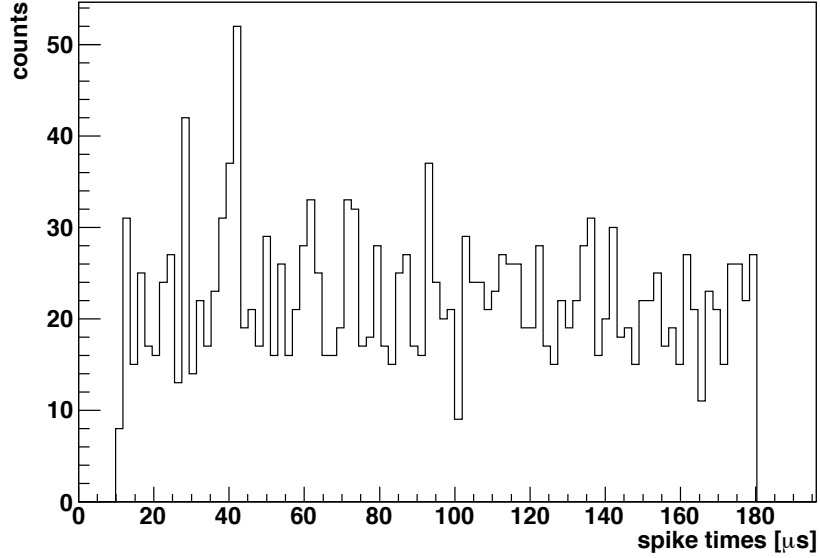


Figure 5.9: Distribution of spike times well beyond the initial scintillation pulse.

pulses and misses the small amplitude scintillation pulses, lowering the spike rate. Furthermore, the 8" PMT discriminator threshold is set quite low, so the measured activity may be higher.

Note that the discrepancy between the pre-trigger region and the post-pulse region may also be due to the pulse timing cut. The first detected pulse is required to be near time zero, so any event with a significant deviation from baseline in the pre-trigger region will show up as a pulse and will be discarded. However, this cut does not play a large role in explaining why the pre-trigger region is quiet or the post-pulse region has an offset. Indeed, the small fraction of events with first pulses found to be before $-0.1 \mu\text{s}$ mostly have small amplitudes (less than half the peak amplitude) and would not have passed the peak amplitude cut anyway.

We find that the post-pulse noise is predominantly explained by subsequent scintillation pulses. After applying the spike cut to the entire post-pulse region for each event, the average waveform is significantly quieter and returns to zero, as shown in Fig. 5.10. Thus, the pre-trigger and post-pulse baseline discrepancy seems mostly due to scintillation photons appearing in the long time window, and a visual check confirms that this is indeed so. (Note that in Sec. 3.6 the spikes were afterpulses.)

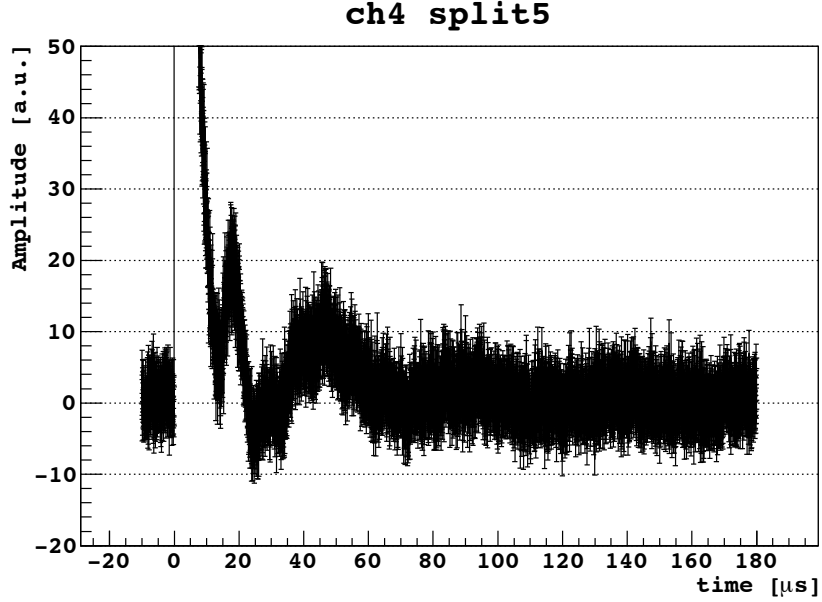


Figure 5.10: The summed waveform for a long time window run with the spike cut.

5.5 Lifetime over time

The slow component lifetime may be affected by impurities in the detector, such as oxygen, which will absorb excitation energy from the argon atoms, thereby reducing the scintillation light yield. The presence of impurities also affects the electron drift lifetime, the time it takes ionization electrons to be drifted to the end of the detector by the applied electric field. The drift lifetime is sensitive to part per billion impurities, while the slow-component lifetime is sensitive to part per million impurities [57]. One of the goals of the DarkSide 10 kg detector was to determine the effectiveness of the LAr purification system. The detector went through several phases of recirculation for purifying the argon, and so it is instructive to compare the behavior of the electron drift lifetimes and slow component lifetimes over long periods of time. Table 5.11 summarizes the behavior of the slow component lifetime over time for a single split in a single channel, and Fig. 5.12 shows the corresponding plot. Figure 5.13 shows the drift lifetimes over time. By eye, the slow-component lifetime is significantly less correlated with recirculation than the drift lifetime.

In analyzing the slow-component lifetime over time, we must also include system-

Date	Run	Slow comp. lifetime
12/12/10	1174*	1.4365 ± 0.0160
12/12/10	1175*	1.4673 ± 0.0069
12/12/10	1176	1.4541 ± 0.0057
12/13/10	1224	1.4608 ± 0.0064
12/16/10	1304	1.4652 ± 0.0059
12/17/10	1323	1.5039 ± 0.0183
12/18/10	1335	1.4844 ± 0.0215
12/21/10	1363*	1.4817 ± 0.0213
12/21/10	1364*	1.4794 ± 0.0105
12/21/10	1365*	1.4833 ± 0.0097
12/21/10	1366*	1.4911 ± 0.0098
12/21/10	1367*	1.4728 ± 0.0094
12/21/10	1368*	1.4609 ± 0.0105
12/21/10	1392	1.4830 ± 0.0062
01/05/11	1439	1.4756 ± 0.0075
01/06/11	1453	1.4312 ± 0.0092
01/07/11	1456	1.4808 ± 0.0100

Figure 5.11: Slow component lifetimes and electron drift lifetimes for a single channel and single split for selected ^{22}Na coincidence runs at null electric field. All lifetimes are shown with statistical errors from the fitter. The systematic error for each lifetime is $0.28 \mu\text{s}$. Runs marked with * were with a collimated source; all other runs were uncollimated.

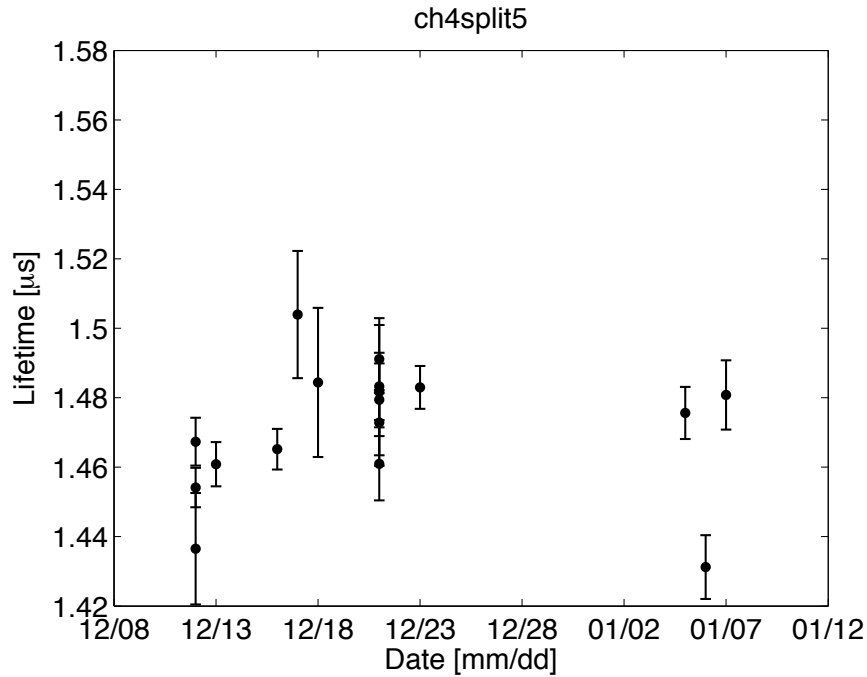


Figure 5.12: The slow-component lifetimes for a single split in a single channel.

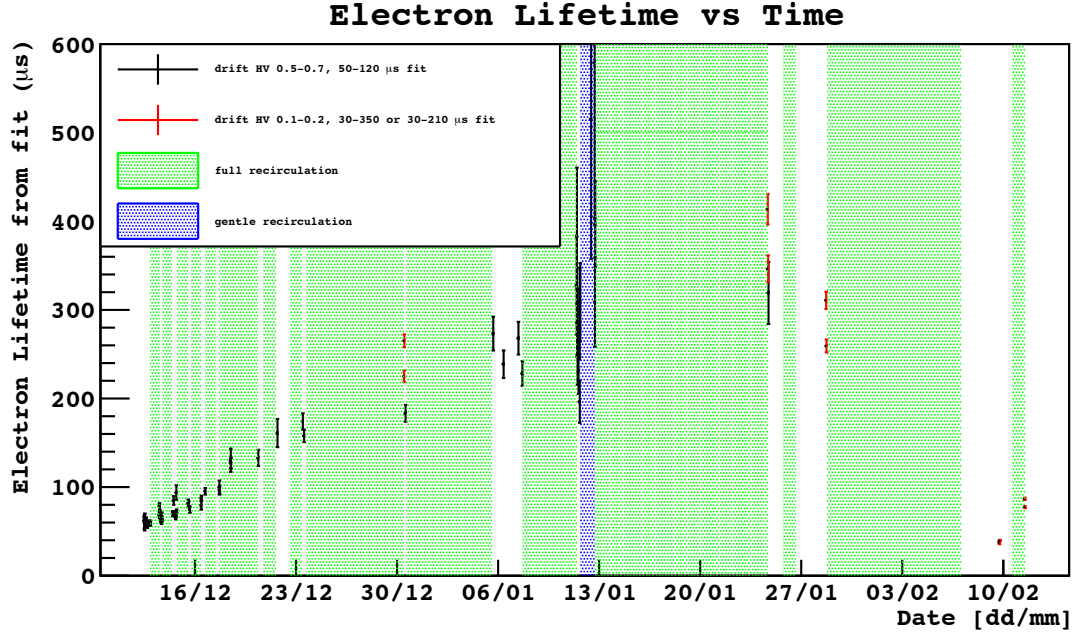


Figure 5.13: The electron drift lifetime over time [53]

atic errors. There is little correlation in the fitted lifetime from channel to channel or from split to split. Therefore we must include their variations in as statistical errors. The rms from split to split is roughly $0.13 \mu\text{s}$ for a single channel, and the rms from channel to channel is roughly $0.15 \mu\text{s}$. Finally we include the variations of the lifetimes due to various fit windows as roughly $0.2 \mu\text{s}$. Adding all the errors in quadrature gives $\sigma_{sys} = 0.28 \mu\text{s}$. This is large compared to the statistical errors of the fits.

Chapter 6

Conclusion

In trying to get a thorough understanding of the slow component in primary scintillation pulses for electronic recoils, we have found multiple anomalous features in the DarkSide 10 kg detector data. We were able to explain the discrepancies between the pre-trigger and post-pulse regions of the long time window runs, but many other features are still unexplained, including the alternating behavior in the baseline data and the double peaks at $\sim 20 \mu\text{s}$ and $\sim 45 \mu\text{s}$ in the summed waveform. While we have a method for handling the alternating behavior, the $\sim 20 \mu\text{s}$ peak may have a strong influence on the fits, so it warrants further study.

Another possible further study is to estimate the slow-component lifetime using a different method. Rather than use a summed waveform, we can measure the individual photon arrival times in each event. These arrival times should be exponentially distributed. This method requires the ability to identify single photoelectrons in the waveform, which may be difficult to do in the first few microseconds where multiple photoelectrons may arrive very near or on top of each other. Thus we can only analyze the later region of the slow component, losing a significant portion of the statistics. On the other hand, this method is less vulnerable to the anomalies found in the summed waveform, especially the alternating baseline and the doubly peaked structure. This timing method would be subject to the same set of cuts as in the summed waveform method, and agreement between the two would greatly enhance our confidence in the calculated lifetime.

Although we cannot give a final absolute value for the slow-component lifetime for electronic recoils, we can still use its calculation as an indicator of the impurity levels in the liquid argon, finding that the lifetime is uncorrelated with recirculation efforts. As the DarkSide experiment begins work on the third iteration of the 10 kg detector and development of the 50 kg detector, the slow component lifetime will have to be visited again.

Appendix A

Sinusoidal Fluctuations

One of the peculiarities of the data coming from the DarkSide dark matter detector is a sinusoidal fluctuation present in all channels. We present here a characterization of the fluctuation.

The fluctuation is most obvious in “quiet” regions of the summed channel waveform. Note that this is a sum over channels—not over events; it is the same sum as described in Sec. ???. Quiet regions are typically the pre-trigger region but also include the post-pulse region of runs long time window, null-field runs such as 1439 and 1456. That the sum appears in the summed channel waveform indicates that the fluctuation is coherent across all the channels. We fit a sinusoidal function to the un-subtracted waveform in these quiet regions, leaving amplitude, period, and vertical offset and horizontal (temporal) offset as free parameters. A sample fit is shown in Fig. [A.1](#)

Fitting each event in run 429, we find that the vertical offset has little variation from event to event and the horizontal offset is approximately uniformly random over a range of 2π , as expected. The fitted period is sharply peaked with mean $4.8 \mu\text{s}$ and $0.8 \mu\text{s}$ spread in its distribution, and the amplitude distribution is peaked at 1.12 counts with spread 0.5 counts. Histograms of the fitted parameters for all events in the run are shown in Fig. [A.2](#). Relative to a typical scintillation pulse, the amplitude of the sinusoidal fluctuation is quite small, but it may have a large effect on the integral of the waveform on which much of the analysis is based. Because the baseline region

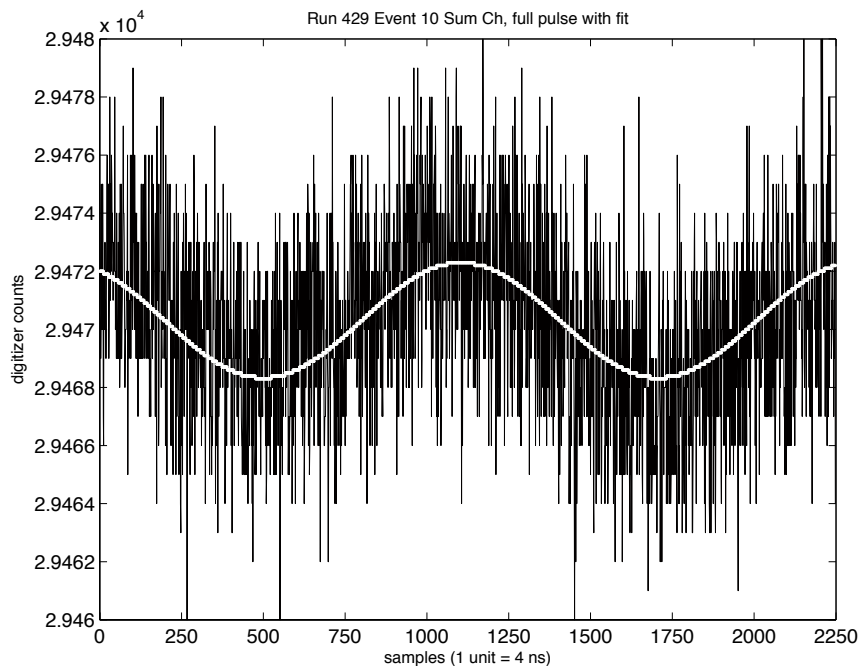


Figure A.1: A sinusoidal fit to the summed channel waveform.

is not typically one period in length, then a simple average calculation for the baseline will likely result in a small deviation from the “true” mean. This causes a runaway in the integrated pulse, making further analysis difficult. Thus, the moving baseline is necessary to tame the integral.

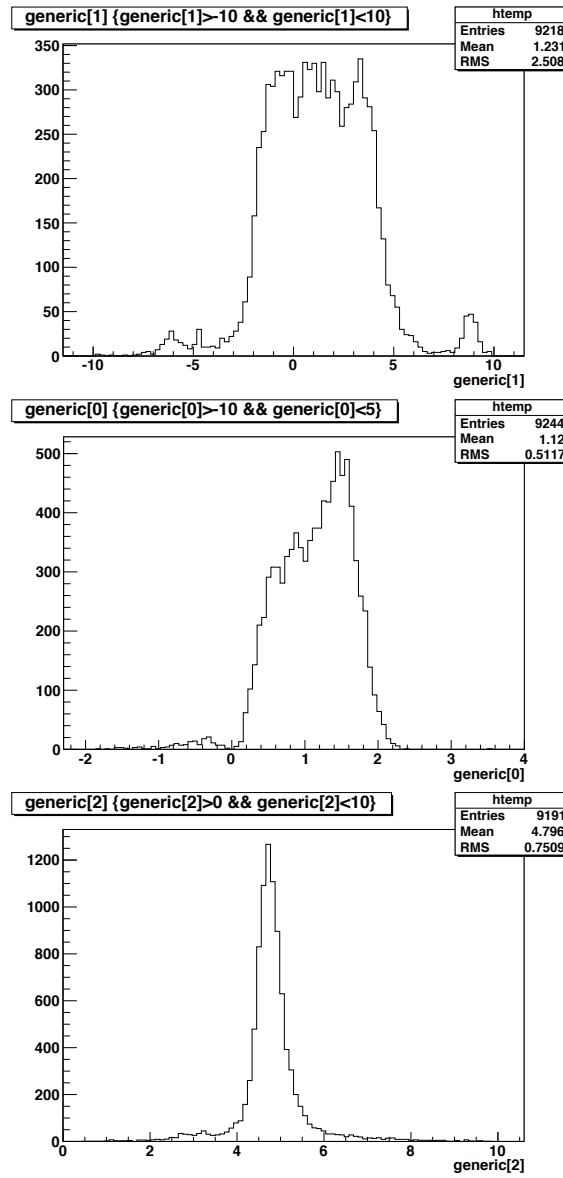


Figure A.2: The distributions of the fitted parameters: period (top), amplitude (middle), and phase (bottom).

References

- [1] N. Bahcall et al., *Science* **284**, 1481 (1999).
- [2] N. Jarosik et al., *Astrophys. J. Suppl.* **192**, 14 (2011).
- [3] K. Nakamura et al., *J. Phys. G* **37**, 075021 (2010), <http://pdg.lbl.gov>.
- [4] G. Bertone, D. Hooper, and J. Silk, *Phys. Rep.* **405**, 279 (2005).
- [5] <http://solarsystem.nasa.gov>.
- [6] G. Clemence, *Rev. Mod. Phys.* **19**, 361 (1947).
- [7] F. Zwicky, *Helv. Phys. Acta* **6**, 110 (1933).
- [8] S. Van den Bergh, *arXiv:9904251* (1999).
- [9] V. Rubin and W. Ford Jr, *Astrophys. J.* **159**, 379 (1970).
- [10] V. Rubin, W. Ford Jr, and N. Thonnard, *Astrophys. J.* **238**, 471 (1980).
- [11] K. Begeman, A. Broeils, and R. Sanders, *MNAS* **249**, 523 (1991).
- [12] D. Clowe, S. Randall, and M. Markevitch, *Nucl. Phys. Proc. Suppl. B* **173**, 28 (2007).
- [13] J. Feng, talk at Princeton Physics Colloquium (2011).
- [14] G. Bertone (ed.), *Particle Dark Matter*, Cambridge University Press, Cambridge (2010).
- [15] S. Asztalos et al., *Annu. Rev. Nucl. Part. Sci.* **56**, 293 (2006).

- [16] A. Szelc, arXiv:1010.3918 (2010).
- [17] G. Bertone, *Nature* **468**, 389 (2010).
- [18] M. Kamionkowski and A. Kinkhabwala, *Phys. Rev. D* **57**, 3256 (1998).
- [19] L. Anchordoqui et al., *Int. J. Mod. Phys. A* **18**, 2229 (2003).
- [20] http://www.nasa.gov/mission_pages/GLAST/news/gr_pulsar.html.
- [21] J. Chang et al., *Nature* **456**, 362 (2008).
- [22] A. Strong and I. Moskalenko, *Adv. Space Res.* **27**, 717 (2001).
- [23] D. Hooper and J. Silk, *Phys. Rev. D* **71**, 083503 (2005).
- [24] http://www.nasa.gov/mission_pages/GLAST/news/fermi-cygnus.html.
- [25] O. Adriani et al., *Nature* **458**, 7238 (2009).
- [26] I. Moskalenko and A. Strong, *Astrophys. J.* **493**, 694 (1998).
- [27] O. Adriani et al., *Phys. Rev. Lett.* **102**, 51101 (2009).
- [28] A. Abdo et al., *Phys. Rev. Lett* **102**, 181101 (2009).
- [29] R. Schnee, arXiv:1101.5205 (2011).
- [30] Z. Ahmed et al., *Science* **327**, 1619 (2010).
- [31] E. Aprile et al., *Noble Gas Detectors*, Wiley-VCH, Weinheim (2006).
- [32] T. Doke et al., *Nucl. Instr. and Meth. A* **269**, 291 (1988).
- [33] E. Aprile et al., *Phys. Rev. Lett.* **105**, 131302 (2010).
- [34] E. Aprile et al., arXiv:1104.2549 (2011).
- [35] <http://deapclean.org/about/>.
- [36] M. Boulay et al., arXiv:0904.2930 (2009).

- [37] D. McKinsey, Nucl. Phys. Proc. Suppl **173**, 152 (2007).
- [38] K. Freese, J. Frieman, and A. Gould, Phys. Rev. D **37**, 3388 (1988).
- [39] A. Drukier, K. Freese, and D. Spergel, Phys. Rev. D **33**, 3495 (1986).
- [40] R. Bernabei et al., Phys. Lett. B **480**(1-2), 23 (2000).
- [41] R. Bernabei et al., Eur. Phys. J. C **56**, 333 (2008).
- [42] H. Loosli, Earth Planet. Sci. Lett. **63**, 51 (1983).
- [43] P. Benetti et al., Nucl. Instr. and Meth. A **574**, 83 (2007).
- [44] D. Acosta-Kane et al., Nucl. Instr. and Meth. A **587**, 46 (2008).
- [45] D. Tadayoshi et al., Jpn. J. Appl. Phys **41**, 1538 (2002).
- [46] W. Lippincott et al., Phys. Rev. C **78**, 035801 (2008).
- [47] P. Benetti et al., Astropart. Phys. **28**, 495 (2008).
- [48] W. R. Leo, *Techniques for Nuclear and Particle Physics Experiments*, Springer-Verlag, New York (1994).
- [49] http://jp.hamamatsu.com/resources/products/etd/eng/html/pmt_001.html.
- [50] K. Arisaka et al., Astropart. Phys. **31**, 63 (2009).
- [51] A. Teymourian et al., arXiv:1103.3689 (2011).
- [52] Image taken from DarkSide internal website (2011).
- [53] (2011), internal memo for the DarkSide experiment.
- [54] Hamamatsu Photonics K.K., *Photomultiplier Tubes* (2006).
- [55] <http://wwwasdoc.web.cern.ch/wwwasdoc/minuit/node4.html>.
- [56] <http://root.cern.ch/drupal/content/users-guide>.

- [57] L. Grandi, WARP: An argon double phase technique for dark matter search, Ph.D. thesis, University of Pavia (2005).